Buyer's Guide: High Availability for Virtual Environments p.71

Windows IT Pro

JUNE 2010 | WINDOWSITPRO.COM | WE'RE IN IT WI

15th ANNIVERSARY Windows IT Pro 15 1995–2010

Group Policy Improvements

in Windows 7 and Windows Server 2008 R2 p. 23

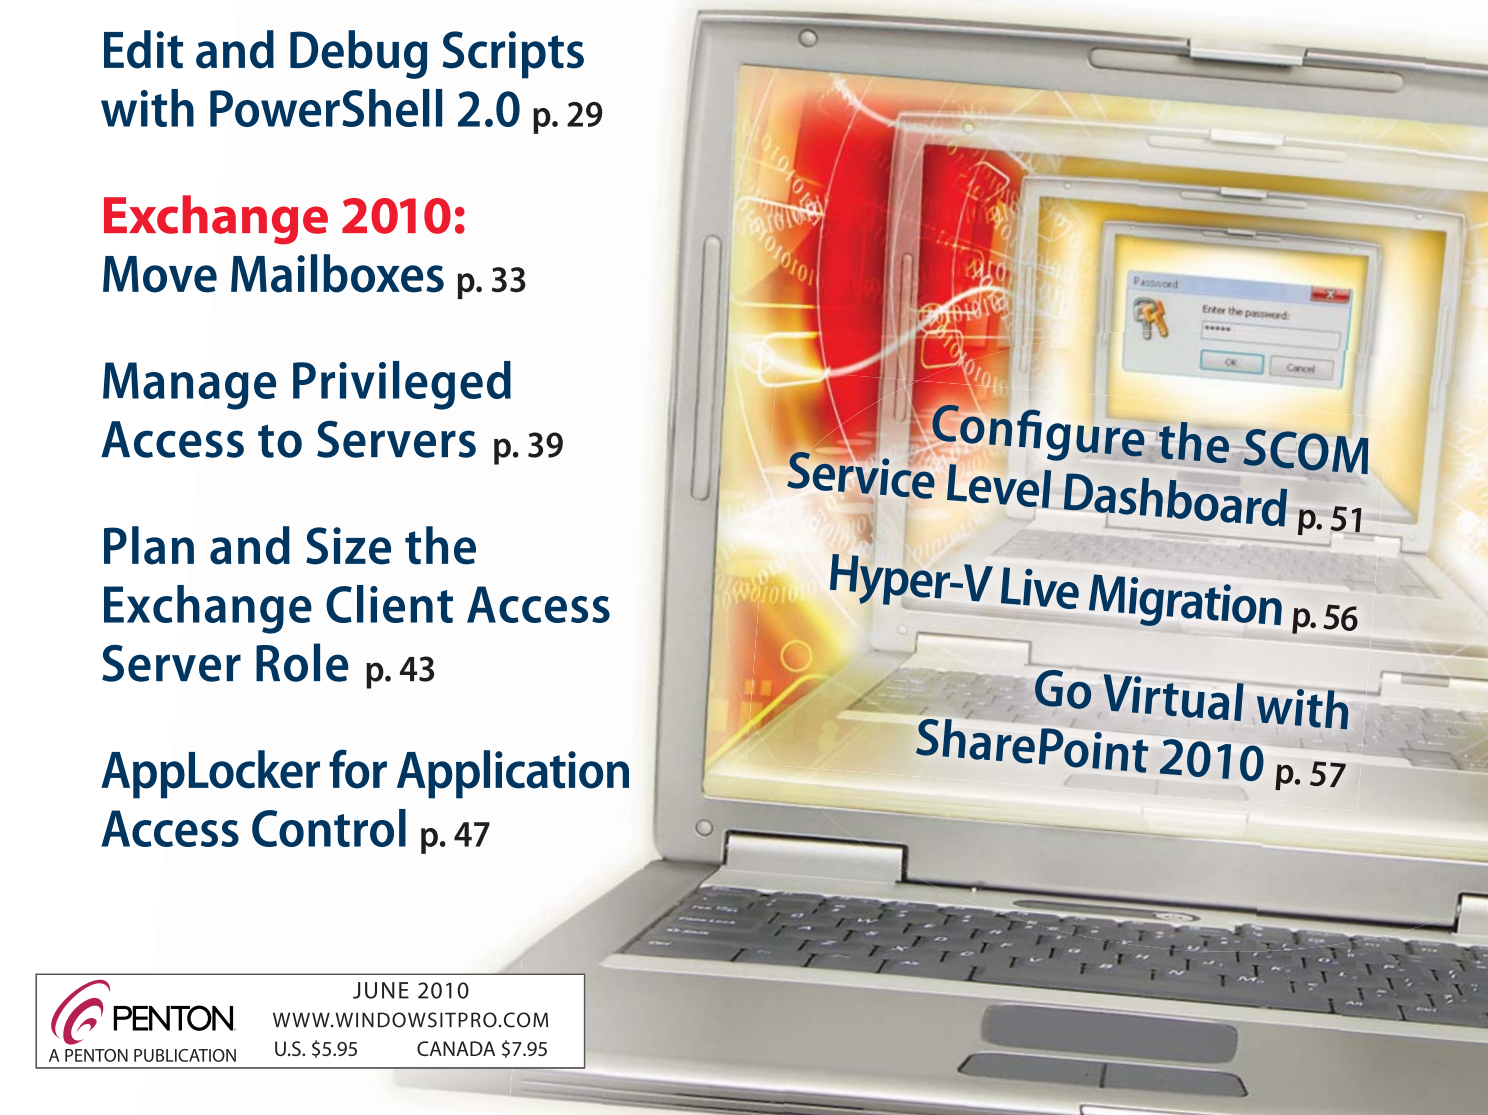

PENTON
A PENTON PUBLICATION

JUNE 2010
WWW.WINDOWSITPRO.COM
U.S. $5.95 CANADA $7.95

Buyer's Guide: High Availability for Virtual Environments p.71

Windows IT Pro

JUNE 2010 | WINDOWSITPRO.COM | WE'RE IN IT WI

15th ANNIVERSARY Windows IT Pro 15 1995–2010

Group Policy Improvements

in Windows 7 and Windows Server 2008 R2 p. 23

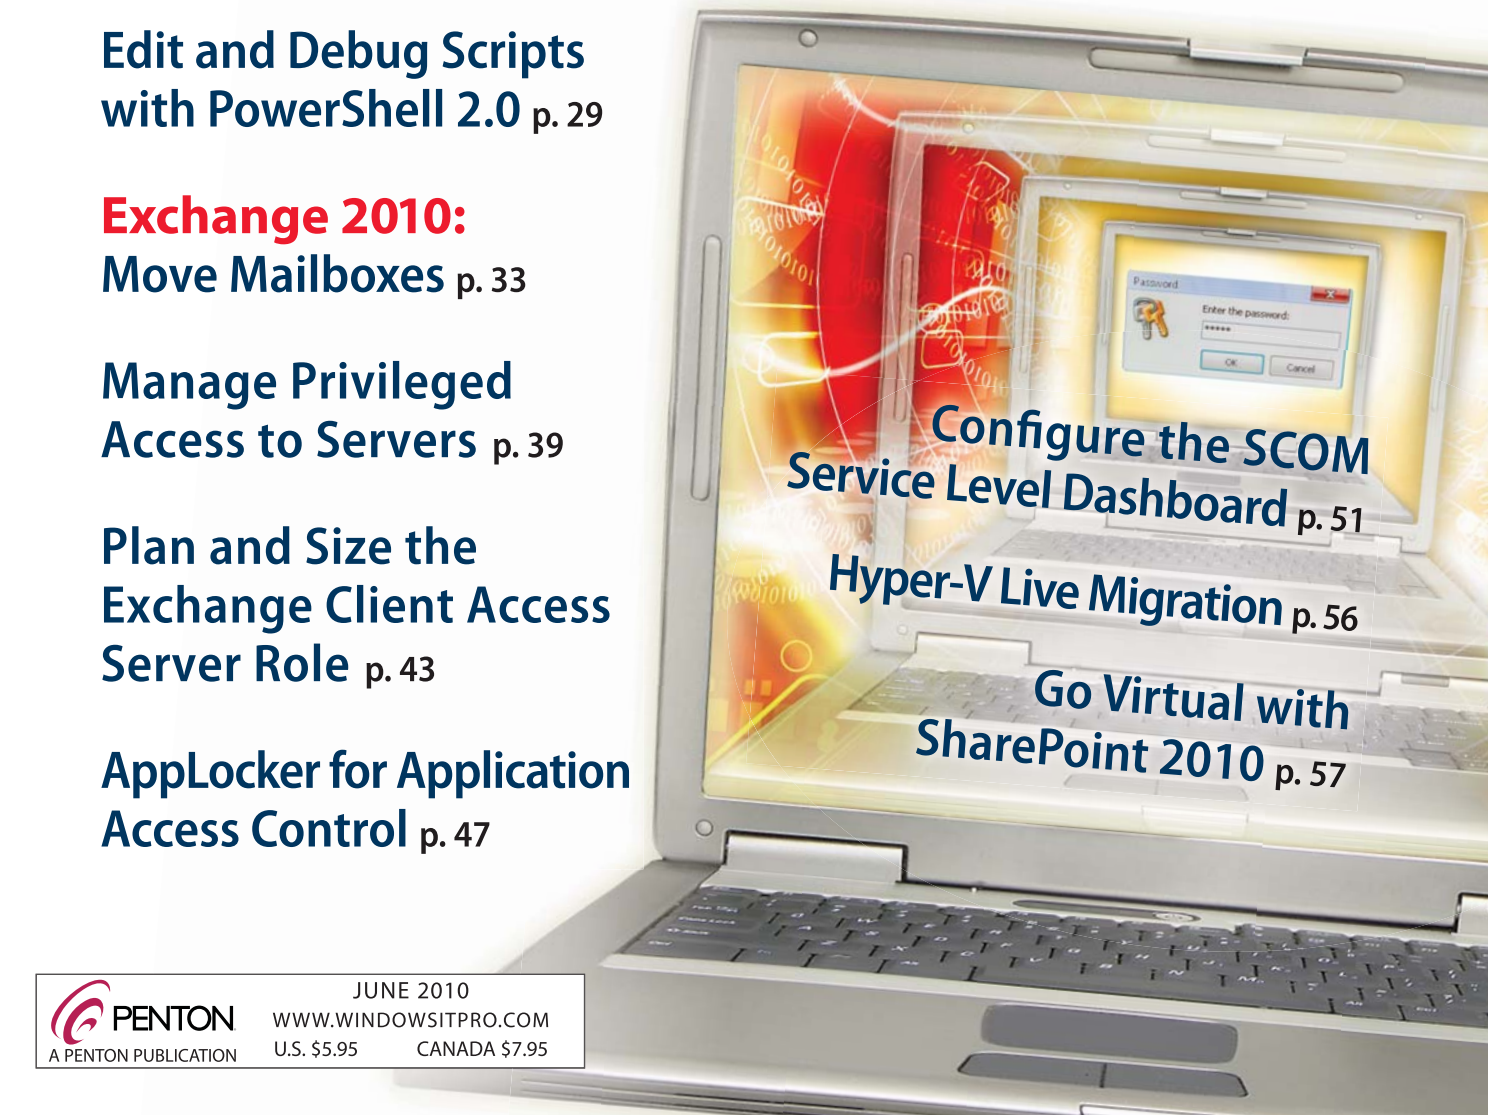

PENTON
A PENTON PUBLICATION

JUNE 2010
WWW.WINDOWSITPRO.COM
U.S. $5.95 CANADA $7.95

TOTAL WINE
DOUBLED IN size
How much new hardware did we buy?
ZERO

Director of Technology
Todd Slan
Total Wine & More

Microsoft

CASE STUDY: Total Wine and More

Total Wine & More is "America's Wine Superstore," operating 54 stores in 11 states and carrying thousands of different wines, spirits and beers. With the company's success and growth has come a proliferation of servers in stores and at headquarters to support new business initiatives—and with those, an increase in technology management work. To relieve the IT staff of the need to manually manage servers, Total Wine took advantage of solutions in the Microsoft® System Center suite of solutions to automate server imaging, application and update deployment, server monitoring, and backup.

With fewer technology management chores, Total Wine can grow its business and roll out new services. Store employees spend more time focusing on customers and accomplishing productive tasks, thanks to more reliable systems, better backup management and fewer local server problems.

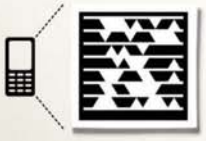

To download the case study, snap this tag or text SAVE to 21710*

Get the free app for your phone at http://gettag.mobi

*Standard messaging and data charges apply.

To read the full case study, visit **itseverybodysbusiness.com/save**

C O N T E N T S

FEATURES

INTERACT

IN EVERY ISSUE

Access articles online at www.windowsitpro.com. Enter the article ID (located at the end of each article) in the InstantDoc ID text box on the home page.

Windows IT Pro
A PENTON PUBLICATION
JUNE_2010
VOLUME_16
NO_6

COLUMNS

.COM

FREE

FOR ONE YEAR*

NO STRINGS ATTACHED!

1&1® INSTANT DOMAIN PACKAGE:

- ✓ **FREE Private Domain Registration**
- ✓ **1&1 Starter WebsiteBuilder**
- ✓ **E-mail Account With 2 GB Mailbox**
- ✓ **24/7 Toll-Free Customer Support**

1&1®

MEMBER OF united internet

Get started today, call 1-877-GO-1AND1

www.1and1.com

*Offer valid as of May 1, 2010 and applies to the Instant Domain Package only. After first year, standard pricing applies. Limit 1 per customer. Visit www.1and1.com for full promotional offer details. Program and pricing specifications and availability subject to change without notice. 1&1 and the 1&1 logo are trademarks of 1&1 Internet AG, all other trademarks are the property of their respective owners.

CONTENTS

PRODUCTS

Windows IT Pro

EDITORIAL

Editorial and Custom Strategy Director
Michele Crockett mcrockett@windowsitpro.com

Executive Editor, IT Group
Amy Eisenberg amy@windowsitpro.com

Technical Director
Michael Otey motey@windowsitpro.com

Senior Technical Analyst
Paul Thurrott news@windowsitpro.com

Custom Group Editorial Director
Dave Bernard dbernard@windowsitpro.com

Web and Developer Strategic Editor
Anne Grubb agrubb@windowsitpro.com

Systems Management
Karen Bemowski kbemowski@windowsitpro.com
Caroline Marwitz cmarwitz@windowsitpro.com
Zac Wiggy zwiggy@windowsitpro.com

Messaging , Mobility, SharePoint, and Office
Brian Keith Winstead bwinstead@windowsitpro.com

Networking and Hardware
Jason Bovberg jbovberg@windowsitpro.com

Security
Lavon Peters lpeters@windowsitpro.com

SQL Server
Megan Bearly Keller mkeller@windowsitpro.com
Sheila Molnar smolnar@windowsitpro.com

Production Editor
Brian Reinholz breinholz@windowsitpro.com

IT Media Group Editors
Linda Harty, Chris Maxcer, Rita-Lyn Sanders

CONTRIBUTORS

SharePoint and Office Community Editor
Dan Holme danh@intelliem.com

Senior Contributing Editors
David Chernicoff david@windowsitpro.com
Mark Joseph Edwards mje@windowsitpro.com
Kathy Ivens kivens@windowsitpro.com
Mark Minasi mark@minasi.com
Paul Robichaux paul@robichaux.net
Mark Russinovich mark@sysinternals.com

Contributing Editors
Alex K. Angelopoulos aka@mvps.org
Sean Deuby sdeuby@windowsitpro.com
Michael Dragone mike@mikerochip.com
Jeff Fellinge jeff@blackstatic.com
Brett Hill brett@iisanswers.com
Darren Mar-Elia dmarelia@windowsitpro.com
Tony Redmond 12knocksinna@gmail.com
Ed Roth eroth@windowsitpro.com
Eric B. Rux ericbrux@whshelp.com
John Savill john@savilltech.com
William Sheldon bsheldon@interknowlogy.com
Randy Franklin Smith rsmith@montereytechgroup.com
Curt Spanburgh cspanburgh@scg.net
Orin Thomas orin@windowsitpro.com
Douglas Toombs help@toombs.us
Ethan Wilansky ewilansky@windowsitpro.com

ART & PRODUCTION

Production Director
Linda Kirchgesler linda@windowsitpro.com

Senior Graphic Designer
Matt Wiebe matt.wiebe@penton.com

ADVERTISING SALES

Publisher
Peg Miller pmiller@windowsitpro.com

Director, International and Agency Services
Don Knox don.knox@penton.com

EMEA Managing Director
Irene Clapham irene.clapham@penton.com

Director of IT Strategy and Partner Alliances
Birdie J. Ghiglione birdie.ghiglione@penton.com 619-442-4064

Online Sales and Marketing Manager
Dina Baird Dina.Baird@penton.com

Key Account Directors
Jeff Carnes jeff.carnes@penton.com 678-455-6146
Chrissy Ferraro christina.ferraro@penton.com 970-203-2883

Account Executives
Barbara Ritter barbara.ritter@penton.com 858-759-3377
Cass Schulz cassandra.schulz@penton.com 858-357-7649

Client Project Managers
Michelle Andrews 970-613-4964
Kim Eck 970-203-2953

Ad Production Supervisor
Glenda Vaught glenda.vaught@penton.com

MARKETING & CIRCULATION

Customer Service service@windowsitpro.com

IT Group Audience Development Director
Marie Evans marie.evans@penton.com

Marketing Director
Sandy Lang sandy.lang@penton.com

CORPORATE

PENTON

Chief Executive Officer
Sharon Rowlands Sharon.Rowlands@penton.com

Chief Financial Officer/Executive Vice President
Jean Clifton jean.clifton@penton.com

TECHNOLOGY GROUP

Senior Vice President, Technology Media Group
Kim Paulsen kpaulsen@windowsitpro.com

WRITING FOR *WINDOWS IT PRO*
Submit queries about topics of importance to Windows managers and systems administrators to articles@windowsitpro.com.

LIST RENTALS
Contact MeritDirect, 333 Westchester Avenue, White Plains, NY or www.meritdirect.com/penton.

REPRINTS
Diane Madzelonka, Diane.madzelonka@penton.com, 216-931-9268, 888-858-8851

.COM FREE FOREVER

Get up to 5 included domains at no additional cost with a 1&1 website plan!*

1&1® HOME PACKAGE

- **2 FREE** Domains (.com, .net, .org, .info or .biz)
- FREE Private Domain Registration
- 150 GB Web Space
- 1&1 WebsiteBuilder
- 1&1 Photo Gallery
- 1&1 Blog
- 24/7 Toll-Free Support

~~$6.99~~

SPECIAL OFFER 3 MONTHS FREE*

1&1® BUSINESS PACKAGE

- **3 FREE** Domains (.com, .net, .org, .info or .biz)
- FREE Private Domain Registration
- 250 GB Web Space
- 25 FTP Accounts
- 50 MySQL® Databases
- 1&1 WebStatistics
- 24/7 Toll-Free Support

~~$9.99~~

SPECIAL OFFER 3 MONTHS FREE*

1&1® DEVELOPER PACKAGE

- **5 FREE** Domains (.com, .net, .org, .info or .biz)
- FREE Private Domain Registration
- 300 GB Web Space
- 50 FTP Accounts
- 100 MySQL® Databases
- PHP 5/PHP 6 (beta) Supported With Zend® Framework
- 24/7 Toll-Free Support

~~$19.99~~

SPECIAL OFFER 3 MONTHS FREE*

1&1®

MEMBER OF united internet

Get started today, call 1-877-GO-1AND1 **www.1and1.com**

*Included domains are free as long as your 1&1 web hosting package is current and in good standing. 3 months free offer valid as of May 1, 2010, a 12 month minimum contract term and a setup fee of $4.99 for the Home Package, and $9.99 for the Business Package and Developer Package apply. Visit www.1and1.com for full promotional offer details. Program and pricing specifications and availability subject to change without notice. 1&1 and the 1&1 logo are trademarks of 1&1 Internet AG, all other trademarks are the property of their respective owners.

Crockett

"Line-of-business upgrades and security products are the most likely to retain funding during tough economic times."

IT Departments Poised to Wield Their Power

IT organizations will drive increased spending in 2010

We're now midway through 2010—the year the economy is supposed to get better—but no one is taking any gains for granted, particularly in the IT world. As the economy rebounds, IT organizations will continue to hold a significant amount of power in determining how precious dollars are spent.

About half of readers responding to a recent *Windows IT Pro* survey reported that the business leaders in their organizations always accepted the recommendations of the IT organization. The other half reported that business leaders sometimes accept those recommendations. Either way, IT professionals will have significant influence this year on how the slowly increasing IT budgets are spent.

Spending patterns during the recession reflect the practical focus that businesses have when financial performance is a company's primary target. More than a third of respondents to our survey reported that short-term cost—the impact on the budget within the next 12 months—was the primary factor in IT buying decisions, along with previous experience with a particular brand. These responses point to companies' risk aversion during hard times.

Although some IT pros responded that they re-evaluate product purchases at least yearly, most products are on a staggered reassessment cycle, with line-of-business (LOB) applications and security applications under review most frequently (at least once a year) and server hardware under review least frequently (about every five years). The review cycle also maps to purchasing habits. As Figure 1 shows, respondents indicated that LOB upgrades and security products are the most likely to retain funding during tough economic times. Some of the least likely products to retain funding are mobile devices, hosted services, and IT management software.

As the economy continues to recover, IT professionals who have helped their companies weather the storm with strategic purchases will be heroes. Keeping the business infrastructure running smoothly and the company poised to seize the coming opportunity requires research skills and experience.

Next month, we'll look at how IT pros are navigating the job market as the economy recovers, and I'll talk with some IT job search companies to provide some insight on how IT pros can best position themselves for career growth. Check out this month's survey (and get a chance to win a $25 gift certificate) at www.windowsitpro.com/go/perspectives.

InstantDoc ID 125100

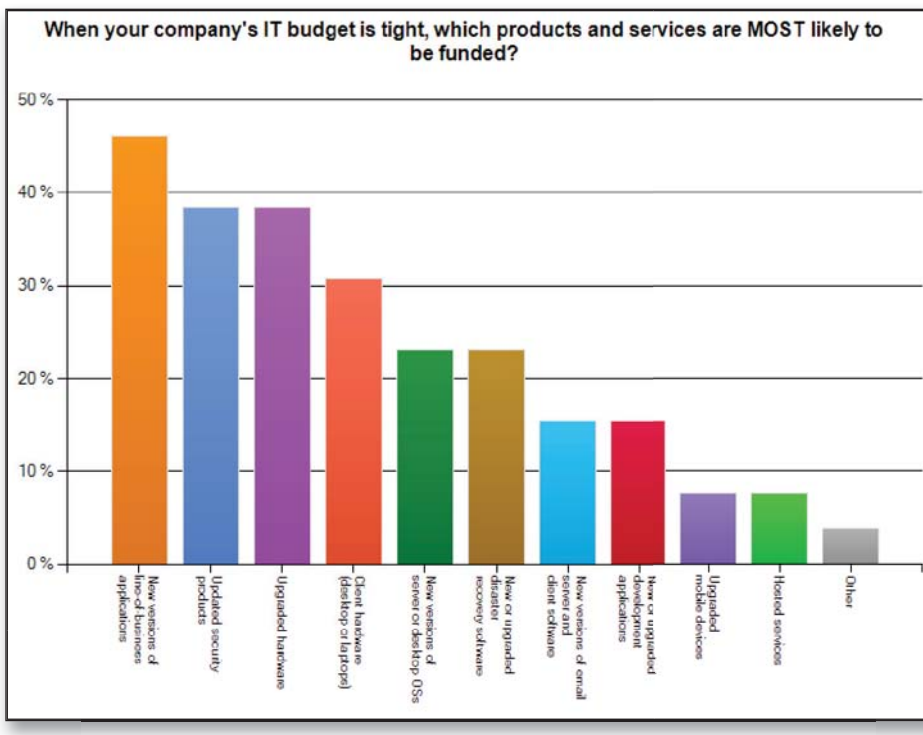

Figure 1: IT products most likely to be funded in tight budget years (April 2010, *Windows IT Pro* Survey)

MICHELE CROCKETT (michele.crockett@penton.com) helped launch *SQL Server Magazine* in 1999, has held various business and editorial roles within Penton Media, and is currently editorial and custom strategy director of *Windows IT Pro, SQL Server Magazine,* and *System iNEWS.*

LETTERS@WINDOWSITPRO.COM

SharePoint Frustrations

In Michael Noel's article, "5 SharePoint Frustrations You Can Overcome" (April 2010, InstantDoc ID 103567), I read something that I don't totally agree with. I've seen this error a few times, so I felt the need to comment.

Regarding Noel's "Content Database Management" section, it's possible to use the Central Administration GUI to determine which content database your new site collection will be created in. The solution isn't elegant, but it works for creating site collections in new and existing content databases and doesn't require that you use Stsadm. Follow these steps:

1. In the Application Management section, access the Content Databases link.
2. In the upper right corner, make sure you're in the correct Web Application. If you aren't, change to the correct one.
3. Click the link for each Content Database you have listed under the Database Name heading.
4. When the Manage Content Database Settings page opens, set the Database Status to Offline, then click OK. The system will take you back to the Manage Content Databases page, where the database you just changed should now appear as Stopped.
5. Click the Add a Content Database link, and create a new content database. Don't make any changes to the Database Status setting. It will appear as Started in the list of Content Databases. You'll also notice that the Current Number of Sites will equal 0.
6. Navigate to the SharePoint Site Management section, and create your new site collection.
7. Return to the Manage Content Databases list, and you'll see that the Content Database you just created now shows that the Current Number of Sites equals 1.

You've just created a new site collection in the content database of your choice. As long as a content database is in the Stopped state, no new site collections can be created within that content database. You can still create sites—or, more accurately, *Webs*—but not site collections.

If you had 15 content databases in the list, and all were in the Stopped state, and you tried to create a new site collection, SharePoint would throw an error message. Pick the content database you want your new site collection to reside in, set the state to Started, and you'll be able to create the site collection in the database you want.

—Jay Simcox

Thanks for writing! I respectfully disagree that my article contains an error. Let me clarify: You're correct that you can use the GUI to accomplish this task. In fact, I mention that in the article. However, you don't need to set the database to a Stopped status to do it. SharePoint uses an algorithm to determine which content database will house a new site collection. This algorithm is based on how much available capacity exists across all content databases. So, the best way to do this is to simply raise the maximum number of sites in the database you want to a very high number, then create the new site collection, which will now go to the database with the most available 'room,' so to speak.

The problem with setting the databases to Stopped is that doing so also prevents site collection admins or anyone in a site from creating sub-sites. In a small environment, that might not be a big deal, but in a large environment it can be a huge concern, as you would effectively prevent everyone from creating new sites or workspaces for the duration of any site-collection creation. Also, when you have a large number of databases, setting them all to Stopped status and changing them back becomes tedious. It's a common misconception that you have to set them all to Stopped to get them to go into the right database; a better option is to raise the maximum number of sites in the content database, as I indicated in the article. Try my solution: You'll find that it works every time, but be sure to set the number very high to ensure that the algorithm puts it in the new database.

I cover this topic in my upcoming SharePoint 2010 Unleashed *book. In addition, you can check out the Shared Points for SharePoint blog (blogs.msdn.com/mcsnoiwb/archive/2007/08/20/how-to-create-site-collection-in-a-specific-content-database.aspx).*

In any case, the point of the article was to show you the silliness of having to go through these types of loopholes to accomplish something relatively simple. Microsoft should just have a mechanism in the GUI that lets you specify a content database.

—Michael Noel

Choosing a Smartphone

I read Brian Winstead's article, "Choosing a Smartphone: The OS" (April 2010, InstantDoc ID 103473). Last year, I bought an AT&T HTC Touch with Windows Mobile 6.1. Stay clear! I think most of my problems were related to Windows Mobile 6.1. I got lucky and discovered a website that showed how to reflash the device, so I upgraded to Windows Mobile 6.5. It's at least bearable now. (I can't tell you how many times I nearly threw the device on the pavement and stomped on it.)

One thing to consider—and the factor that ultimately drove me toward this particular phone—is that I didn't want to pay

Windows IT Pro welcomes feedback about the magazine. Send comments to letters@windowsitpro.com, and include your full name, email address, and daytime phone number. We edit all letters and replies for style, length, and clarity.

Inventor of the PC

Thank you, Paul Thurrott, for mentioning Ed Roberts in WinInfo Daily UPDATE (April 2, 2010). I'm sorry about Ed's passing and will gratefully remember him. MITS gave me my start in the computer industry in 1976, when I first met Bill Gates, Paul Allen, and many other great young computer enthusiasts. Thank you for helping people understand that a small company in Albuquerque—and not IBM or Apple—really created the industry that supports the world today.

—Randall Huddleston,
President, Censerve Consulting

$50 per month for a data plan. I just don't want to browse the web on a 3" screen. Most phones (e.g., from Apple and Google) require you to get the data plan. The guys at T-Mobile, Verizon, and AT&T simply wouldn't sell their nice phones without one.

—Douglas Nebeker

Thanks for sharing your personal experience with the HTC Touch and Windows Mobile 6.1. Because of HTC's great reputation, I looked seriously at the HTC Imagio on Windows Mobile 6.5 but ultimately heard too many stories of freeze-ups and poor performance. As for the data plan problem, you're right that it should be a major consideration. However, since my employer is covering the plan, I didn't have control over that detail. I knew from the beginning that I needed something on Verizon, and the data plan was included. My situation probably echoes that of many people shopping for business-use phones.

I've since written a few other articles on this topic, and in fact I've chosen—and begun using—the Motorola Droid. Also, I'll be writing a more in-depth Market Watch article for the July issue, about the various mobile OSs and devices and what IT shops need to do to support them.

- "Choosing a Smartphone: The Hardware," www.windowsitpro.com/go/winsteadhardware
- "Choosing a Smartphone: The Features," www.windowsitpro.com/go/winsteadfeatures
- "Choosing a Smartphone: My Choice," www.windowsitpro.com/go/winsteadmychoice
- "The Wonderful World of Droid," www.windowsitpro.com/go/winsteadwonderfuldroid

—Brian Keith Winstead

What's So Bad About Vista?

I've been reading Paul Thurrott's WinInfo News for years and generally agree with his comments. But I have to wonder: Exactly what's so bad about Windows Vista? The supposed faults have totally escaped me from its release date to the debut of Window 7.

Before switching to Windows 7 (which I've been very happy with), I ran Vista for years, and as a Visual C++ developer, I have a greater understanding of an OS than the average user. I never once had to re-install one of my Vista computers; they never misbehaved. In fact, I couldn't really fault Vista at all throughout the years I used it across all machines.

I struggle to understand the constant press about how bad Vista was and how great Windows 7 is. Actually, Windows 7 is about the same as Vista. OK, it's faster to boot up and shut down, and the UI is slightly nicer in some respects. But the Windows 7 experience is hardly different from that of Vista.

—Nic Wilson

Instant Poll Results: Windows 7 Migration

When do you plan to migrate your business to Windows 7?

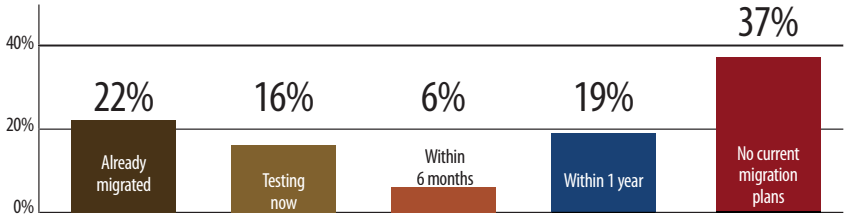

Source: *Windows IT Pro* Instant Poll, www.windowsitpro.com, April 2010.

Disabled Virtualization on 64-Bit Processors

Regarding John Savill's FAQ (InstantDoc ID 104690) about 64-bit processors that can't be used for virtualization, there have been a [very] few cases in which laptop manufacturers have disabled virtualization capabilities to increase battery life of the laptop. (I don't know whether the change was permanent or could be changed in the BIOS.)

If that's the case, I believe the manufacturer is shortchanging the client. If I bought a laptop with virtualization capabilities, it should be my decision to enable or disable them. I should at least be notified that it's been disabled and given the option to enable it. Imagine buying a car with the air conditioning disabled because it uses extra gas!

—Ed Braiter

InstantDoc ID 125123

Thurrott

"SCE 2010 hits all the midmarket high points, providing automation and central management for the busy IT generalists that keep this market running."

What You Need to Know About New Service Packs, System Center Essentials, DPM 2010, and Apple iPad

This month, I've got news, tips, and opinions about yet another great set of technologies and products. And go figure, most of them are enterprise products, for a change. Here's what you need to know.

SP1 for Windows 7 and Windows Server 2008 R2

Microsoft announced something I've been dying to discuss: details about the first service pack for Windows 7 and Windows Server 2008 R2. (You might recall that both products are on the same code base and thus are serviced by the same service packs.) Last month, I noted that Windows 7 SP1 would mostly aggregate the software updates that appeared during its first year. However, Windows Server 2008 R2 SP1 will include some major new features:

Dynamic memory. Anyone comparing Microsoft's Hyper-V virtualization platform and the VMware stack will conclude that Microsoft's solution is evolving rapidly but lacking in a few key areas. Well, one of those remaining areas will be addressed with dynamic memory support in SP1. Dynamic memory makes it possible to pool the available memory on the host server and dynamically distribute it to virtual workloads as needed. In other words, you can dynamically allocate RAM to virtual machines (VMs), on the fly, without needing to shut them down first.

RemoteFX. This functionality—which came via Microsoft's purchase of Calista—is now finally available via SP1. It enhances the display experience during RDP sessions, supporting Windows Aero, Microsoft Silverlight, and Adobe Flash user experiences, and 3D graphics. (Citrix is also partnering with Microsoft to integrate RemoteFX technologies into HDX for XenDesktop.)

Microsoft hasn't announced the timing of SP1, but the plan is to ship it one year after Windows 7 and Windows Server 2008 R2. And as you know, those products dropped in October 2009.

Windows HPC Server 2008 R2

Speaking of Windows Server, Microsoft is working to finalize Windows HPC Server 2008 R2. A public Beta 2 version is available now, and Microsoft hopes to ship the final version by the end of 2010. HPC Server 2008, formerly marketed as Windows Compute Cluster Edition, is Microsoft's entry in the very high-end high-performance computing (HPC) server market. HPC Server 2008 R2 is aimed at the most scalable server systems and provides parallel computing capabilities. Most exciting is its ability to run Microsoft Excel calculations in parallel on a cluster, providing the absolute best performance for scientists, engineers, and analysts.

Exchange 2010 SP1

Service packs are in the air, apparently. Despite Exchange 2010's recent launch, Microsoft is already at work on SP1. Yes, Exchange 2010 SP1 will aggregate the product and security fixes that shipped since the product launched. But it will also include new capabilities, including better mail archiving, Outlook Web App (formerly Outlook Web Access) improvements, and better management functionality. Let's look at these closely.

Mail archiving. Although Exchange 2010 includes integrated mail archiving capabilities, this feature is enhanced in SP1. It will offer the option to provision individual users' personal archives to a different mailbox database other than their primary mailbox, import historic email from client-side PST files, and provide controls so that admins can delegate access to a user's Personal Archive. A new multi-mailbox search feature will make it easier to find email for legal or regulatory reasons.

Outlook Web App. Microsoft's web-based Exchange client is being updated again with a visual refresh, better performance through pre-fetching and asynchronous delete, the ability to mark as read and categorize operations, and public shared access to calendars. Some features from previous OWA versions are returning, too, like UI themes, and the ability for users to move the reading pane.

Management. With SP1, the Exchange Management Console (EMC) and Exchange Control Panel (ECP) are being enhanced to provide access to new management tasks, many of which relate to Retention Policy Tags that can help automate how email is deleted and archived.

Other new features are coming in Exchange 2010 SP1 as well, and it's starting to shape up as a major release. Microsoft expects to deliver it by the end of 2010.

System Center Essentials 2010

Although Microsoft cancelled its midmarket-focused Essential Business Server product line—I was told that customers appreciated the integrated management but not the requirement that they buy three to four servers at a time—it hasn't lost focus on this small but

important part of its customer base. Indeed, with Microsoft's midmarket customers returning to mainstream Windows Servers and other Microsoft servers, they need good management tools. So, not surprisingly, the software giant is offering a new version of its midmarket management solution, called System Center Essentials 2010.

SCE 2010 hits all the midmarket high points, providing automation and central management for the busy IT generalists that keep this market running—and if you're familiar with the System Center family of products, you'll understand why this is such a big deal. (Microsoft defines midmarket as mid-sized companies with up to 50 servers and 500 clients.) By the time you read this, the final release of SCE 2010 will have hit the streets on June 1, 2010. Here's what you get.

Unified management of physical machines and VMs. SCE has always offered a unified management experience, but in SCE 2010 this has been enhanced with virtualization capabilities from Microsoft System Center Virtual Machine Manager (VMM), including the ability to manage physical and virtual servers and clients side by side, provisioning abilities, snapshots (which Microsoft calls Checkpoints), physical-to-virtual conversion, conversion from VMware to Hyper-V, machine migration, physical resource optimization, and more. This is the biggest improvement in SCE 2010 by far and represents an aggressive move into the virtualization space for this market segment. (Previously, Microsoft offered a bundle of the separate SCE 2007 SP1 and VMM 2008 tools.)

Monitoring and reporting. SCE's monitoring capabilities help overworked IT staff in midsized businesses move from a reactive stance to being more proactive, while the reporting capabilities provide a running, high-level view of the health of the environment. The nicest thing about the UI, however, is that all of the status messages are hot links, providing not just information but actionable targets that can help admins fix problems. Embedded videos are attached to some links, helping admins learn on the go without leaving the SCE management environment.

Software deployment and update management. SCE 2010 builds off the deployment and update management capabilities from previous versions. There are many enhancements, and Microsoft offers a third-party catalog so that deployment and update functionality isn't limited to its software. SCE 2010 offers more granular deployment and update management than before, so you can do things such as target machines that meet certain criteria (like 32-bit versions of Windows XP running particular languages). This is a nice solution for those seeking to deploy complex applications such as Office 2010.

Software and hardware inventory. As with its predecessors, SCE 2010 provides a complete asset inventory system for software and hardware in your environment. This helps with licensing compliance but also provides guidance when you want to upgrade to new OSs or applications.

Additionally, while SCE 2007 shipped with a core set of management packs, SCE 2010 streamlines this and offers a more intelligent approach where only those management packs that apply to the software in your environment are surfaced. And unlike before, you can see these relevant new packs in the centralized management console rather than having to check the Microsoft website.

System Center Data Protection Manager 2010

System Center Data Protection Manager (DPM) 2010 is the first major update to Microsoft's data protection product since DPM 2007 SP1, and it expands on that product's core functionality around protecting file shares, Exchange, Microsoft SQL Server, Microsoft SharePoint, and Hyper-V. Oriented to protecting only Microsoft workloads, DPM 2010 integrates with the core data storage technologies in each (such as Shadow Copies in Windows Server), and provides near-term restoration from disk and long-term restoration from tape. New 2010 features include these:

Cloud restore capabilities. Building on an exclusive offer that first appeared in DPM 2007 SP1, DPM 2010 supports replicating data to an Iron Mountain vault, extending the previous disk and tape restoration capabilities.

Client support. For the first time, DPM also supports client-based data protection. And for laptops and other clients that might be disconnected from the corporate network, DPM 2010 can create local shadow copies to a reserved area of the disk, then protect the data on the server when a connection is established to the network. DPM 2010 supports 1,000 clients per server and works with both Windows 7 and Windows Vista.

DPM 2010 will ship alongside SCE 2010 on June 1. If you're interested in deploying both, Microsoft is offering a new Essentials Plus License that combines a single client management license for each into a single, less expensive package. Otherwise, pricing and licensing for SCE 2010 and DPM 2010 haven't changed.

Apple iPad

And now for a change of pace, let's consider Apple's iPad. Launched with much fanfare and hype, the iPad is basically a very large iPod touch and not a miniature, tablet-based Mac. This design decision has positives and negatives, but I think it was the right choice, and as the iPad app ecosystem improves, and Apple lowers prices, it could become an interesting consumer option.

Today, however, I can't recommend the iPad to most people, and it's no value to most businesses. Stay tuned to this space, however: Apple will improve the iPad, and many competitors—including HP, which is prepping a Windows 7-based Slate PC—will pop up over time, providing highly portable machines that aren't as limited as Apple's offering. I know it's not acceptable to be anything but overly enthusiastic about an Apple product, but that's where I'm at right now.

Windows 7 Enterprise Trial

Finally, here's a tip: If you're interested in evaluating Windows 7 but haven't done so, Microsoft has extended the availability of the trial version of Windows 7 Enterprise through the end of 2010. (Previously, it was a limited time offer, though it's unclear how the company planned to measure when it had "run out" of downloadable copies.) Learn more at the Microsoft Springboard website at tinyurl.com/win7eval.

InstantDoc ID 125039

PAUL THURROTT (thurrott@windowsitpro.com) is the news editor for *Windows IT Pro*. He writes a weekly editorial for *Windows IT Pro UPDATE* (www.windowsitpro.com/email) and a daily Windows news and information newsletter called *WinInfo Daily UPDATE* (www.wininformant.com).

Minasi

"A VHD that has been selected, attached, partitioned, formatted, and given a letter is almost indistinguishable from a physical hard disk."

Diskpart Exerts VHD Control

Manipulate virtual hard disks as if they're actual physical disks

In "Diskpart Goes Virtual" (InstantDoc ID 103685), I showed you how to use Diskpart's Create Vdisk command to build a Virtual Hard Disk (VHD) file. A VHD is a file that lets you easily package and transport data that can be mounted to act as a hard disk or, of course, used as a virtual disk with Microsoft's virtualization tools. There's more to using a VHD than creating it, however, so this month I'll show you a few more VHD-related commands.

Last month, you created a 200MB VHD named *e:\test.vhd*. Now that you have this VHD, you'd like to interact with it (e.g., examine its contents, put files in it), so you'll need to use a couple more commands in Diskpart: Select Vdisk and Attach Vdisk. In our example, the Select Vdisk command and its results look like

```
DISKPART> select vdisk file=e:\test.vhd
DiskPart successfully selected the virtual disk file.
```

Select Vdisk doesn't offer any options except *noerr*, an option available on most Diskpart commands that simply instructs Diskpart to keep running a script despite errors.

You still can't look inside the VHD or create files in it, though. Once you've created and selected the VHD, you can use Attach Vdisk to instruct Windows to treat the VHD as if it were an actual physical disk. Its simplest syntax looks like

```
attach vdisk
```

Note that you don't need *file=<vhdfilename>* parameters. Diskpart simply examines the virtual disk that you've just designated with the Select Vdisk command and instructs Windows that you've installed a new hard disk. You can now use commands such as Create Partition, Assign Letter, and Format to prepare the virtual disk for normal disk operations. Once your VHD has been selected, attached, partitioned, formatted, and given a letter, it's almost indistinguishable from a physical hard disk.

When you've finished creating files and installing software on your VHD, you can instruct Windows to stop treating the VHD as an actual hard disk by typing two commands:

```
select vdisk file=vhdfilename
detach vdisk
```

We discussed the Create Vdisk command last month, but you probably won't use that command nearly as often as you'll use the Select Vdisk/Attach Vdisk and Select Vdisk/Detach Vdisk combinations. For example, if you have three already-created VHDs on your hard drive, you could type three Select Vdisk/Attach Vdisk command pairs and end up simultaneously working with three VHDs attached as three "imaginary" hard disks.

Suppose you sit down at a system that you know nothing about, and you see a surprising number of lettered volumes. You want to know how to determine which, if any, of the system's hard drives are VHDs. If you're working from a system with a full GUI, the easiest way to find the VHDs is to look at the drive icons in Logical Disk Manager (diskmgmt.msc); LDM colors VHDs cyan instead of the normal gray. If you don't have a GUI to work with, you can use Diskpart's List Vdisk command. On a system with two attached virtual disks, I get the output that Figure 1 shows.

Recall the List Disk command, which lists all the drives on a system. On a system with VHD support, List Disk lists only those drives that seem somewhat local. I say "somewhat local" because although mapped network drives aren't listed, iSCSI LUNs are listed—and so are attached VHDs. Unfortunately, List Disk doesn't offer much in the way of clues to a disk's true nature, so you can't quickly see which are iSCSI or VHD. You *can* always determine a disk's nature by first selecting it (Select Disk *disknumber*), then typing *Detail Disk*. It will then report a disk type (e.g., iSCSI, Virtual).

```
DISKPART> list vdisk

  VDisk ###  Disk ###  State                Type       File
  ---------  --------  -------------------  ---------  ----
  VDisk 0    Disk 2    Attached not open    Fixed       c:\test2.vhd
* VDisk 1    Disk 1    Attached not open    Fixed       c:\test.vhd
```

Figure 1: List Vdisk results

But that's a pain, which is why List Vdisk is so likeable. The second column of List Vdisk's results—Disk ###—refers to the disk number shown in the List Disk output. Nice!

Oh, and one more thing about attached VHDs in Windows: The attachment doesn't survive a reboot, so if you're going to need a VHD attached all the time, consider cooking up a Diskpart script for your logon or startup script. Happy attaching!

InstantDoc ID 125054

MARK MINASI (www.minasi.com/gethelp) is a senior contributing editor for *Windows IT Pro*, an MCSE, and the author of 25 books, including *Mastering Windows Server 2008 R2 (Sybex)*.

Otey

"Although some of these shortcuts provide a quicker way to perform the same actions you're able to perform using the mouse, others have no graphical equivalent."

Windows 7 Keyboard Shortcuts

Use the keyboard to switch between applications, launch Windows Explorer, lock your desktop, and access other handy features

The debate has long raged whether keyboard shortcuts or mouse clicks are a more efficient way to accomplish tasks, and of course the answer will usually come down to personal preference. However, if you're using a laptop keyboard frequently, you might find some keyboard shortcuts essential. In this Top Ten column, I'll share my favorite Windows 7 keyboard shortcuts. These shortcuts let you work quickly and efficiently with your Windows desktop. Although some of these shortcuts provide a quicker way to perform the same actions you're able to perform using the mouse, others have no graphical equivalent. And while some of these shortcuts have been available in Windows Vista and Windows XP, others are brand new to Windows 7.

10 **Windows key+Spacebar**—The Windows key (Win)+Spacebar keystroke combination gives you a quick way to display your desktop. All open Windows are made transparent, and you can see the desktop underneath. However, you can't interact with anything on the desktop. When you release the Windows key, the open windows become visible again.

9 **Win+Home**—This keyboard shortcut provides a handy way to quickly clean up your workspace. Pressing Win+Home minimizes all open windows with the exception of the currently active window.

8 **Win+Number**—The Windows key plus a number provides a shortcut to opening items in the Windows 7 taskbar. Press the Windows key together with the number that corresponds to the taskbar item you want to open, counting left to right. For instance, Win+2 opens the second taskbar item.

7 **Win+(+ or -)**—This keyboard shortcut can be handy during presentations because it lets you magnify the screen. Pressing the Windows key and the plus sign (+) makes the entire screen larger with each iteration. Conversely, pressing the Windows key and the minus sign (-) makes the screen display progressively smaller.

6 **Win+T**—This keyboard shortcut is handy if you want to quickly scroll through the different items in the Windows 7 taskbar. Pressing Win+T cycles through the taskbar items, displaying the thumbnail image for each item. If you want to change which item you're currently working in, you can press Enter when an item is highlighted to switch to that item.

5 **Win+Left arrow and Win+Right arrow**—Although some of these shortcuts have been in previous versions of Windows, these two are totally new to Windows 7. They let you take advantage of the UI's side-by-side docking feature. Pressing the Windows key and either the Right or Left arrow key docks the current Window to the side of the desktop that corresponds with the direction of the arrow.

4 **Win+L**—This keyboard combination is super handy for quickly locking your desktop. Pressing Win+L locks the desktop and displays the Windows logon screen. To unlock the desktop, you need to enter your Windows password.

3 **Win+E**—Even faster than launching Windows Explorer by right clicking the Start button, pressing Win+E launches Windows Explorer, starting with the Computer view. From there, it's easy to navigate through the system's drives shown in the right portion of the Windows Explorer window to find whatever you're looking for.

2 **Win+P**—If you've ever needed to give presentations on multiple brands of laptops, you know how annoying it can be to search each different type of laptop for its external projector or monitor hot key. Windows 7 fixes that problem. Pressing Win+P displays the Presentation Display Mode window, which lets you toggle your laptop's presentation mode between Computer Only, Duplicate, Extended, or Projector Only.

1 **Win+X**—If you're running Windows 7 on a laptop, this keyboard shortcut is for you. Pressing Win+X displays the Windows Mobility Center, which lets you control a number of system settings, including the audio level, the power scheme, wireless networking, external displays, and the external projector. The Windows Mobility Center is also often customized by each OEM.

InstantDoc ID 125032

MICHAEL OTEY (motey@windowsitpro.com) is technical director for *Windows IT Pro* and *SQL Server Magazine* and author of *Microsoft SQL Server 2008 New Features* (Osborne/McGraw-Hill).

Stock

"In this two-part article series, I'll discuss the tools used by the Microsoft Support Team to troubleshoot kernel memory corruption."

Troubleshooting Kernel Memory Corruption

To solve memory corruption "crimes," first learn how to spot telltale clues left in pool memory

As an IT administrator, you've no doubt been the victim of the infamous drive-by blue screen crash. It typically happens at the most inconvenient times, often resulting in service interruptions and work stoppage causing monetary loss. Working on the Microsoft Windows escalation team, I'm tasked with debugging these crashes daily and reporting a timely action plan to prevent such problems from happening again. The work is a lot like the popular crime show, *CSI*, where crime scene investigators use low-level techniques to flesh out the guilty party. In some of my debugging investigations, the DNA left behind is as microscopic as a single bit flipped to a 1 when the code was expecting a 0. In a high percentage of the cases, the offender is long gone except for the memory corruption that remains, making it extremely difficult to find the misbehaving driver. How do we weed out culprits that cause memory corruption?

This article kicks off a two-part article series, in which I'll discuss the tools used by the Microsoft Support Team to troubleshoot kernel memory corruption typically caused by a buggy driver that doesn't kindly leave behind its calling card. Because of the complexity of the memory manager, it may not be immediately clear to you why the support team prescribes these tools. This article and the next one should clear up any confusion and will be especially helpful if you're responsible for reporting an action plan back to your management team. I should also mention that although there are many other reasons for bug-check crashes, this article will focus primarily on some of the tools used by Microsoft Support to diagnose crashes caused by kernel memory corruption.

Pool Memory Architecture

Before I delve into the tools, I'll provide a short primer on high-level pool memory architecture, which may help you better understand when I explain the tools in more detail. Most of the memory that drivers use is allocated from the system-provided pools, called *paged pool* and *nonpaged pool*. There are exceptions to this rule, but a discussion of them is beyond the scope of this article. As the names imply, the memory used by the nonpaged pool is always guaranteed to be resident in physical memory, whereas portions of the paged pool can be swapped out at any given point in time.

To better understand the output from our diagnostic tools, it's important to note that a pool is divided into smaller divisional units called *pages*, either small or large pages. A small page is 4KB. A large page is either 2MB on the x64 platform or 4MB on x86 systems. However, if the Physical Address Extension (PAE) option is enabled on x86, the large page size is reduced to 2MB. Both page types have their advantages, but for the rest of this article we'll assume I'm talking about the smaller 4KB pages.

A page is further subdivided into the actual allocations made by the various drivers asking for memory. Figure 1 shows a random page of nonpaged memory that I dumped, which provides an example of a 4KB memory page.

I used the !pool command to display the memory in the Microsoft debugger (windbg.exe) that's available with the Microsoft Debugging Tools from the Microsoft download site at www.microsoft.com/whdc/devtools/debugging/default.mspx. (For more information about using the Microsoft debugger, see "Administrators' Intro to Debugging," June 2009, www.windowsitpro.com/go/windbg.) Windbg is the primary debugging tool used by the Microsoft Global Escalation Services team. (In future articles I plan to

```
*fffffadcda813000 size:  100 previous size:    0  (Allocated) *Mdl
 fffffadcda813100 size:   10 previous size:  100  (Free)       ....
 fffffadcda813110 size:   50 previous size:   10  (Allocated)  FLli
 fffffadcda813160 size:   20 previous size:   50  (Free)       Even
 fffffadcda813180 size:   40 previous size:   20  (Allocated)  VadS
 fffffadcda8131c0 size:   60 previous size:   40  (Allocated)  Even (Protected)
 fffffadcda813220 size:   80 previous size:   60  (Allocated)  SeTd
 fffffadcda8132a0 size:  150 previous size:   80  (Allocated)  Time (Protected)
 fffffadcda8133f0 size:   60 previous size:  150  (Allocated)  Even (Protected)
 fffffadcda813450 size:   b0 previous size:   60  (Allocated)  MmCa
 fffffadcda813500 size:  1d0 previous size:   b0  (Allocated)  CcSc
 fffffadcda8136d0 size:   50 previous size:  1d0  (Allocated)  Afdh (Protected)
 fffffadcda813720 size:   80 previous size:   50  (Allocated)  Ntfr
 fffffadcda8137a0 size:   60 previous size:   80  (Allocated)  Even (Protected)
 fffffadcda813800 size:   60 previous size:   60  (Allocated)  Even (Protected)
 fffffadcda813860 size:   60 previous size:   60  (Allocated)  Even (Protected)
 fffffadcda8138c0 size:   60 previous size:   60  (Allocated)  Even (Protected)
 fffffadcda813920 size:  2f0 previous size:   60  (Free)       AfdP
 fffffadcda813c10 size:   40 previous size:  2f0  (Allocated)  VadS
 fffffadcda813c50 size:   40 previous size:   40  (Allocated)  VadS
 fffffadcda813c90 size:  370 previous size:   40  (Allocated)  Irp
```

Figure 1: 4KB memory page example

discuss the different debugging techniques used with Windbg for diagnosing these types of crashes.)

This output represents a single 4KB page of nonpaged memory. Each row is a block of memory either allocated to or freed from the page. Most of the blocks in this example are allocated, and you can sort of determine the owner of the block by the tag listed to the right of each allocation.

When a driver makes a call to allocate pool memory, it passes the size of the requested block along with a four-character identifier called a *tag*. The tag and the size of the allocation requested are maintained in a bookkeeping structure called the *pool header*, which is parked at the top of each allocation. This structure also maintains the current size and previous size block, which is used by the memory manager to easily traverse the contents of the page for maintenance tasks such as coalescing adjacent freed blocks into one large freed block. The area following the pool header is arguably the most important area from the perspective of the drivers allocating the memory, as this is the actual storage area for the memory page's data.

Debugging a blue-screen crash is a lot like the popular crime show, *CSI*, where crime scene investigators use low-level techniques to flesh out the guilty party. Often the offender is long gone except for the memory corruption that remains.

The Crime Scene: Dusting for Driver-Bug Fingerprints

Building on the context of how the memory manager organizes pool memory, let's discuss what happens when software bugs creep into the environment and cause those unpleasant blue screen crashes. One of the most common driver bugs is when the driver writes beyond its allocation, spilling data into the next allocation and overwriting data it doesn't own. As mentioned earlier, the pool header precedes the actual data area in each allocation, so when the driver writes into the next allocation, the pool header for the next allocation becomes corrupted. If the memory manager later attempted to read from the corrupt pool header, the system would possibly crash with a *Bug Check 0x19: BAD_POOL_HEADER* or *Bug Check 0xC2: BAD_POOL_CALLER*. The parameters of these bug checks are documented fairly well on MSDN. A look at the first two parameters usually indicates the state of the corruption; however, it doesn't point you to the buggy driver.

Let's extend the scenario before we investigate the tools we'll use to weed out the misbehaving drivers. Using the output in Figure 1, assume the driver using the VadS pool at virtual address fffffadcda813c50 wrote beyond the pool header of the Irp allocation at fffffadcda813c90 and continued writing into the data section. Now we have a scenario where not only the pool header is corrupted, but the driver data, too. If the driver owning the Irp pool used the corrupted data, there is a high likelihood of system instability. Even worse, the owner of the Irp pool would appear to be the guilty party because that driver might be on the stack if the system crashed.

The bug-check code would vary depending on how the owner of the Irp pool used the corrupt data. It may manifest as a STOP 0x0000001e if the driver of the Irp pool attempted to de-reference the corrupt value as a pointer and the value was inaccessible. And what if the address was accessible and the driver wrote data to this random address? Now the corruption runs into another pool or perhaps a critical kernel structure. I point all this out to illustrate how extremely difficult it can be to trace back to a driver when there is no distinct path back to the guilty party.

Learning Path

More articles on Windows debugging:

"Administrators' Intro to Debugging"
www.windowsitpro.com/article/performance/administrators-intro-to-debugging.aspx

"Bit Flips: Was That a Zero or a One?"
www.windowsitpro.com/article/administration-tools2/bit-flips-was-that-a-zero-or-a-one-.aspx

"Find the Binary File for Any WMI Class"
www.windowsitpro.com/article/windows-management-instrumentation-wmi/find-the-binary-file-for-any-wmi-class.aspx

"Find the Source of an Error Message"
www.windowsitpro.com/article/utilities/find-the-source-of-an-error-message.aspx

"Further Adventures in Debugging"
www.windowsitpro.com/article/utilities/further-adventures-in-debugging.aspx

"Q. My machine is crashing and is showing a blue screen. How can I find out what's causing the crash?"
www.windowsitpro.com/article/utilities/q-my-machine-is-crashing-and-is-showing-a-blue-screen-how-can-i-find-out-what-s-causing-the-crash-.aspx

While investigating this scenario, we could easily point the finger at the VadS pool, making this an open-and-shut case. But consider the case where the VadS pool has been freed and another driver allocated the block in its location, only after VadS had corrupted the Irp pool. Now the VadS owner is long gone, making this a cold case file. Enter Special Pool.

The Chase Continues

Introduced in Windows NT SP4, Special Pool was created to catch drivers corrupting memory in real time by allocating guard pages around the allocation. The idea is to catch a driver writing beyond its allocation by forcing it to write into a guard page, causing the system to crash immediately with the culprit on top of the stack. It's the smoking gun approach. In my next article, we'll take a deep look at the Special Pool mechanism and discuss how it's implemented.

InstantDoc ID 125143

RON STOCK (ronsto@microsoft.com) is an escalation engineer for Microsoft's Global Escalation Services team. He specializes in advanced Windows debugging and performance-related issues. For information about Windows debugging, visit his team's blog at blogs.msdn.com/ntdebugging.

- Active Directory Authentication Services
- Recycle Bin
- DC Demotion

READER TO READER

Dogsled Replication Protocol Provides Active Directory Authentication Services When a WAN Isn't Available

I work for a small IT contractor for the federal government. My primary role is to provide Active Directory (AD) design and support for a national federal agency with about 500 discrete locations across the United States. One day, while sitting in my office a few years ago, I got a call from one of my remote site administrators in Northern Alaska, who had an interesting problem. The administrative office where he was working had just opened a new research facility on top of a glacier. It would be open for eight months. During this period, anywhere from 20 to 50 scientists would be working there full time for one or more weeks at a time. The problem was that the scientists were visiting from various locations and had some specific work requirements. They all needed to share data on the two Windows Server 2003 servers in the research facility, and they all needed to share data on the Windows 2003 servers at their home offices. We were asked to come up with a way for them to use the same logon information for the servers in their home offices and in the glacier research facility.

James R. Day

This sounded like a job for AD, so we quickly wrote up a plan. We would join their servers to our existing AD forest, create a *glacier researcher* security group, and add the scientists to it. We would then make this group a member of the local groups at the research facility so we could assign appropriate permissions on the server shares. We could even use a Group Policy Object (GPO) to apply specific research-site settings by taking advantage of the GPO's loopback processing and logon-script capabilities. We only needed to decide if the scientists would log on to a local domain controller (DC) or if the network connection was stable and fast enough to allow them to log on to a remote DC.

Although the scientists were all highly trusted and no financial or personally identifiable information (PII) data was accessible through AD at the time, using a local DC wasn't recommended due to a lack of physical security. So, I called the remote site administrator back and asked him about the WAN.

His answer: "There is none."

"What do you mean there is none?" I asked.

His response was rather surprising. It seems that the research facility was about 200 miles away from the administrative office and didn't have any type of infrastructure going into it. Power was provided by a diesel generator at the site, there were no phone lines, there was no water, and worst of all, there was no Internet or network connectivity of any type.

They had discussed using a wireless bridge, but it turned out that running repeaters up and over two mountains without power wasn't feasible. They had also discussed using satellite connections, but the research facility was located too far north and the providers at that time were below the horizon. Their only contact with the rest of civilization was a weekly supply run, when one member of the team would make a 200-mile trek into town for groceries and other necessities every Thursday.

At this point, we had a solution for their logon issues but no way to extend the domain up there. I jokingly suggested that we create a new replication protocol called "replication over dogsled" (to go with replication over IP and replication over SMTP). We all laughed, but then I thought about it—it would work.

The solution I came up with would provide domain logon services and anything else we needed to the scientists. It also would provide replication of a timely enough nature so that a password change made either at a scientist's home office or the glacier site would be available at the other location. Group changes would also be available at all sites, and we would avoid the problems of lingering objects on DCs (which happens when a DC doesn't replicate during the forest-set tombstone lifetime.) Any object that's deleted on a DC would get stored for the tombstone lifetime, then permanently deleted. If a DC didn't replicate with the rest of the forest for longer than the tombstone lifetime, any deleted object that's present in the replicating DC would be seen as a new object and would be brought back from the grave. My colleagues and I refer to these DCs that don't replicate past the tombstone lifetime as "tombstoned DCs."

The Proposed Solution and Benefits

The solution was simple enough so that we wouldn't need any type of IT support beyond setting it up. The solution would work like this:

We would build a new site (called the *glacier site*) for the glacier research facility and a site-link that connected it to the administrative office 200 miles away. We would place two DCs in the glacier site. The primary DC would be built on standard server hardware that plugged into an uninterruptible power supply (UPS) system connected to the generator. The second DC would be built on a laptop with two network cards: the laptop's internal card and a PC card. The internal card would be configured

Tell the IT community about the free tools you use, your solutions to problems, or the discoveries you've made. Email your contributions to r2r@windowsitpro.com.

If we print your submission, you'll get $100.

Submissions and listings are available online at www.windowsitpro.com. Enter the InstantDoc ID in the InstantDoc ID text box.

with the IP address for the glacier site's network. The PC card would be configured with the IP address for the administrative office's network. The laptop would also be set as the preferred bridgehead for the glacier site. This would ensure that it would be the only DC at the glacier site that would replicate with the administrative office's server.

The site-link replication schedule would be set so that replication would only be allowed during a certain period. During most of the week, the laptop would remain at the glacier site and replicate with the primary DC there. On Thursday, it would travel to the administrative office on the weekly supply run. The laptop DC would stay at that site until Sunday morning. During this time, it would replicate with the administrative office's server. Later on Sunday when the laptop DC returned to the glacier site, it would begin replicating with the glacier site's primary DC. The laptop's power management settings would be configured so that it would go into standby mode when unplugged, then wake up when plugged in. This would prevent the laptop from shutting down unexpectedly and eliminate the need for somebody to log on and power it off.

Every Friday while the laptop was connected to the administrative office's network, I would verify replication by running the command

```
repadmin /showrepl
```

at the main office where I work. By checking the laptop's last replication time with both the administrative office's DC and the glacier site's primary DC, I'd be able to verify that replication was occurring in both places. If either of the replication times was off by more than a few days, we could conclude that either the wrong network card got plugged in or the server didn't get powered on, and work with the IT people from there.

I presented this idea to the rest of my team and to the remote site administrator, pointing out its benefits:

- The scientists could use the same logon credentials from their home offices and from the research facility.
- Permissions for the research facility's servers could be managed via AD groups, greatly simplifying resource management. The administrative office and the scientists' home offices were already members of our AD domain, so these servers already had the necessary permissions. Adding the glacier site's servers to the domain would require a bit of front-end work, but once that would be completed, we could use the existing AD groups to manage permissions at all the sites where the scientists worked.
- Changes made at the administrative office or anywhere else in the AD domain would replicate to the glacier site within one week, keeping things more or less current. This would include changes to the scientists' user accounts as well as required policy or configuration changes (e.g., locking the computer after inactivity) made by security or desktop administrators at other sites.
- Changes made at the glacier site would replicate back to the rest of the domain within one week, keeping the entire domain relatively up-to-date.
- Scientists who changed their password (or any other AD attribute) at either the research facility or their home office would find that change at the other location the following week. Because the scientists were either at the research facility or their home location for a week at a time, the one-week lag time between updates would be acceptable.
- The research facility's primary DC and laptop DC wouldn't become tombstoned due to replication failures for longer than the tombstone lifetime of 60 days. The primary DC was kept current with changes by the laptop DC from Sunday night until Thursday, and the laptop DC was current with the rest of the domain by the weekend.
- Site-based GPO settings could be enforced at the glacier site.

Everybody agreed to this solution, which was dubbed the "dogsled replication protocol." We built it over the following week. The only change made was the decision to avoid using actual dogsleds, replacing them with a 4X4 SUV that was able to travel the road up to the top of the glacier.

We were able to maintain this solution for five months until the SUV was involved in an accident that destroyed the laptop DC. At that point, the decision was made to close the facility for the winter.

The Possible Applications and Risks

The dogsled replication protocol can be adopted for use in other scenarios. A couple of scenarios that come to mind are as follows:

- An office needs to be set up for a short time and WAN links aren't available or aren't feasible. This could include offices for summer camps, retreats, or emergency work locations.
- An office is suffering from a long-term WAN outage caused by damaged wireless transmitters, satellite receivers, or cable connections. The damage might be caused by the environment (e.g., storms, earthquakes) or even people. In one case, I used the dogsled replication protocol for a few months at a site where a mile of copper wire was stolen and the site wasn't able to replace it for several months.

There are some security risks associated with this solution. The biggest risk concerns the physical security of both the onsite DC and the laptop DC. The onsite DC can be somewhat physically secured, depending on the specific situation. However, if someone gains physical access to the onsite DC, he or she has logical access to your entire AD database. This is the same security risk that Microsoft documents for any branch-office scenario.

The laptop DC poses a much higher security risk. The people transporting this DC must be trusted individuals, as they'll have unsupervised physical access to the server for prolonged periods. There's also the risk of theft or damage to the laptop DC while it's in transit.

When we implemented this solution, we felt that these risks were limited because of population factors (nobody lived near the glacier site or on the road to it) and the type of data stored in AD. AD didn't touch any of our financial or messaging systems at that time and didn't contain any PII data. It provided authentication and access to research-based file shares only.

With the security improvements in Windows Server 2008, there are a couple of actions you can take to mitigate the security risks associated with the onsite DC. You can set up the onsite DC as a read-only domain controller (RODC), caching only the passwords of the branch office's users. This significantly decreases the risk of the onsite DC being compromised. The downside is

that changes (including password changes) can't be made on the RODC.

You can't set up the laptop DC as an RODC because it must be a writeable DC—that's the only way to replicate changes between two locations. However, you can use Windows BitLocker Drive Encryption to protect the laptop DC.

For example, had BitLocker been available when we used the laptop DC in Alaska, we could've set up BitLocker as follows. We could've stored the startup encryption key on two USB flash drives. One USB flash drive could've remained in a locked drawer or safe at the glacier site, while the other one could've remained in a locked drawer or safe at the administrative office. We could've then scheduled a task to shutdown the laptop DC at 5 A.M. every Thursday and at 5 A.M. every Sunday so that it would be powered off during transport. When the laptop DC arrived at the research facility or administrative office, someone could've started it by inserting the USB flash drive, plugging in the network cable, and pressing the power button. Using BitLocker in this way would have protected the laptop's information from being stolen, compromised, or used while in transit.

A Viable Solution

By using Windows 2003 and some creative replication settings, we were able to provide AD authentication services to a large number of people over a long period of time in a location where a WAN connection wasn't available. Although there were some security risks, these risks can now be lessened with Server 2008 and BitLocker. Thus, if you find yourself in a situation where WAN isn't available or feasible, you might consider trying the dogsled replication protocol.

—James R. Day,
senior system engineer, NuAxis

InstantDoc ID 125056

Working with Recycled Files in VBScript and PowerShell

Back in the 1990s, my first major Windows scripting project was automating a system cleanup and migration. One deceptively simple task was emptying the Recycle Bin. I finally ended up having the script delete the folder during the cleanup process.

Alex K. Angelopoulos

Since then, tools like the Shell32.dll component have simplified working with special folders such as the Recycle Bin. For example, the VBScript script in Listing 1, EmptyRecycleBin.vbs, performs the task that I was trying to perform during the migration years ago—it empties the recycle bin without prompting. Like any Recycle Bin operation these days, this script has its limits: It only empties the current user's Recycle Bin, and if some of the items need elevated permissions to be deleted, you need to run the script with elevated permissions.

While EmptyRecycleBin.vbs is useful for deleting all the items in the Recycle Bin, it doesn't let you delete individual items. That's where the PowerShell script Get-Recycled.ps1 comes in handy.

Get-Recycled.ps1 enumerates all items in the Recycle Bin. While enumerating these items, the script adds the following information:

- The date the item was deleted (DeletionTime and DeletionTimeUtc)
- The directory in which the item was originally located (OriginalParent)
- The item's original name (OriginalName)
- The item's complete original path (OriginalFullName)
- How long the item has been in the Recycle Bin (Age)

You can then use this information to delete or perform another type of operation on specific Recycle Bin items. For example, Listing 2 shows some PowerShell commands that use this information to do the following: delete files that have been in the Recycle Bin more than 10 days, get the original location of all the files deleted in the last hour, and restore all the .zip files in the Recycle Bin.

Listing 1: EmptyRecycleBin.vbs

```
Option Explicit
Dim sa, fso, item, items
Set sa = _
  CreateObject("Shell.Application")
Set fso = _
  CreateObject("Scripting.FileSystemObject")
Set items = sa.Namespace(10).Items()
On Error Resume Next
For Each item in items
  If fso.FileExists(item.Path) Then
    fso.DeleteFile item.Path, True
  Else
    fso.DeleteFolder item.Path, True
  End If
Next
```

Listing 2: Sample Commands That Use Get-Recycled.ps1

```
# Delete items that have been in the
# Recycle Bin more than 10 days.
Get-Recycled | ?{$_.Age.TotalDays -gt 10} |
  Remove-Item


# Get the original location of all items
# deleted in the last hour.
Get-Recycled | ?{$_.Age.TotalHours -le 1} |
  %{$_.OriginalFullName}


# Restore all .zip files without
# overwriting newer files with the
# same names.
Get-Recycled | ?{$_.Extension -eq ".zip"} |
  %{Move-Item $_ $_.OriginalFullName}


# Restore all .zip files, overwriting
# newer files if necessary.
Get-Recycled | ?{$_.Extension -eq ".zip"} |
  %{Move-Item $_ $_.OriginalFullName -Force}
```

You can download these sample commands, Get-Recycled.ps1, and EmptyRecycleBin.vbs from the *Windows IT Pro* website. Go to www.windowsitpro.com, enter 125117 in the InstantDoc ID box, click Go, then click the *Download the Code Here* button.

—Alex K. Angelopoulos, IT consultant

InstantDoc ID 125117

Windows Server 2008 Allows DC Demotion in Safe Mode

To forcibly demote a Windows Server 2003 or later domain controller (DC), you can run the command

```
dcpromo /forceremoval
```

Murat Yildirimoglu

at a command prompt. However, in Windows 2003, this command doesn't work if the DC was booted in safe mode. Thus, if you have a Windows 2003 DC that won't boot normally, you can't demote it.

In Windows Server 2008 and later, this is no longer a problem. You can use the dcpromo /forceremoval command while in safe mode, allowing you to demote a problem server. It's a welcome change in Server 2008.

—Murat Yildirimoglu, MCSE and MCT

InstantDoc ID 125055

3 TOP **TECH INITIATIVES** TARGETED BY **FRAGMENTATION**

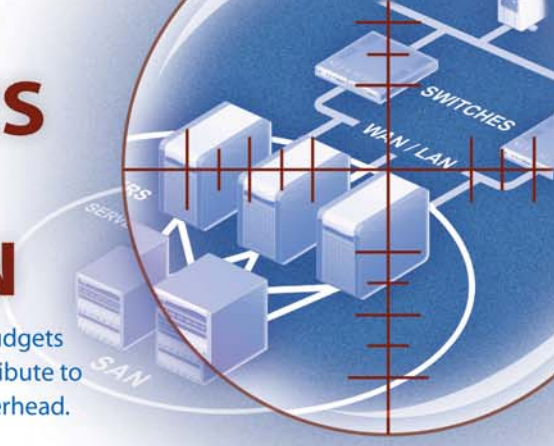

As CIOs and IT managers gear up to meet the challenges of stringent budgets and new tech initiatives, how they handle file fragmentation will contribute to the difference between cost-effective consolidation and increased overhead.

Virtualization

Efficiency vs. "fragmentation on top of fragmentation"

The hard disk is the slowest component of a system's throughput. File fragmentation only makes the bottleneck worse. In the case of virtualization, the disk must do far more; it must support numerous simultaneous operating systems and a greatly compounded rate of fragmentation both on the logical disk and the virtual disks.

These virtual disk files fragment just as any other file can, resulting in what amounts to a "logically" fragmented virtual hard disk, which still has typical file fragmentation contained within it. In other words, virtualization brings about a "fragmentation on top of fragmentation" that can quickly cripple system speed and negate the efficiency virtualization is designed to deliver.

Data Storage Management on SAN Devices

Is fragmentation still really an issue?

A storage area network (SAN) provides the ability to make remote disks appear to be local. SAN storage virtualization involves the creation of a usually very large, logical pool of data. Via software, that pool appears to be physically located all on one server. In actuality, that data may be located across hundreds of physical disks spread across dozens of servers.

The local disk file system does not know of and cannot control the physical distribution or location in a virtualized storage environment. As a result of fragmentation, NTFS has to make multiple requests regardless of the physical or virtualized storage environment.

SANs cannot address file system level fragmentation and neither can proprietary architectures or data retrieval technologies. The overhead on the operating system is heavily impacted by fragmentation. *Local disk file defragmentation is vital.*

The Standard Operating Environment

Lowering network operating costs with efficiencies of scale

There are multiple dynamics that make up overall network efficiency but because file fragmentation is created at the operating system level regardless of how much free space is on the disk, its negative effect on the network is one of the most basic issues to resolve. When not effectively addressed, fragmentation creates a perfect storm of network issues including:

- Slow read/write times
- Slow backups and higher failure rates
- Database lockups
- Shorter productive disk life
- Spiraling energy costs
- Slow boot time
- Increased Help Desk traffic
- Higher re-imaging costs

Resolving fragmentation at base image level would clearly make sweeping improvements to a network, lowering the cost of ownership with the least amount of effort.

The Economics of Fragmentation Prevention

Diskeeper® 2010 technology and the system performance paradigm

Eliminating fragmentation as a performance issue has four basic goals: the reestablishing of optimum performance, reliability, longevity and energy efficiency in every system on a network. Only Diskeeper 2010 includes the innovative functionality to achieve this:

- It prevents up to 85% of all fragmentation before it occurs
- It eliminates any remaining fragmentation in real time
- It quickly handles even the largest mission-critical enterprise servers
- It is completely automatic and invisible
- It includes a centralized graphical administration console scalable to any size

In reality, since every system fragments, any global solution must meet stringent requirements or its operational overhead will negate gains. Diskeeper 2010, with an edition for every Windows® system from laptops to the largest mission critical enterprise servers, is the only solution that increases performance and lowers total cost of ownership at the same time.

Diskeeper 2010

The only way to prevent fragmentation *before* it happens™

CERTIFIED FOR Windows Server 2008

Compatible with Windows 7

Special Offer

Try Diskeeper 2010 FREE for 45 days!

Download at www.diskeeper.com/specialtrial

(Note: Special 45-day trialware is only available at the above link)

Volume licensing and Government / Education discounts are available from your favorite reseller or call 800 829-6468

- System Center Configuration Manager
- FTP
- Kerberos
- BitLocker
- Windows 7

ANSWERS TO YOUR QUESTIONS

Q: By default, Windows displays a user's account information when the user locks a desktop. Is there some way to change this behavior and hide account information from the Computer Locked dialog box?

A: Yes, this behavior can be changed using a registry hack. In a Windows domain environment, you can also use a Group Policy Object (GPO) setting.

The GPO setting is called Interactive Logon: Display User Information when the session is locked and is located in the following GPO container: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options. When you enable this setting, you can set one of these three options:

- User display name, domain and user names—registry value 1
- User display name only—registry value 2
- Do not display user information—registry value 3

The corresponding registry key is called DontDisplayLockedUserId (REG_DWORD) and is located at HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System. To hide the account information from the Computer Locked dialog box, set this registry key to a value of 3.

A side effect of hiding the account information from the Computer Locked dialog box is that when you try to log in to a locked machine, Windows won't show the name of the user who's currently logged on in the logon dialog box. To unlock your logon session, you must type your password and retype your account name.

—Jan De Clercq

InstantDoc ID 104687

Q: I have a large number of BMC files on my computer. What are they?

A: When you use RDP to communicate with hosts, there are times when the same bitmap is repeatedly required on the client. Instead of repeatedly sending the same image to be rendered, the server sends the bitmap once and tells the client to cache the bitmap in an uncompressed form. This cached copy goes in RAM and, optionally, persists (is stored) on disk so it can be used in future sessions. This cache consists of several large BMC files that are stored under %USERPROFILE%\AppData\Local\Microsoft\Terminal Server Client\Cache. You could delete them, but unless you have a real disk space problem, I'd advise against it.

You can move this cache if you want by navigating to HKEY_CURRENT_USER\Software\Microsoft\Terminal Server Client using RegEdit and creating a new String value named BitmapPersistCacheLocation. Set this value to the location where the bitmap cache files should be stored.

—John Savill

InstantDoc ID 125067


Jan De Clercq | jan.declercq@hp.com
John Savill | jsavill@windowsitpro.com

Q: How can I stop Microsoft System Center Configuration Manager (SCCM) site systems from storing data on certain partitions?

A: By default, SCCM will use NTFS drives that have space on them. However, if you have drives you don't want SCCM to use, you can stop it by placing an empty file named no_sms_on_drive.sms on the root of the partition that shouldn't have SCCM content.

— John Savill

InstantDoc ID 104696

Q: How can I block FTP users from taking down my Windows FTP server by filling up its disk space?

A: Since Windows 2000, Microsoft has included a disk quota utility in the Windows OS that you can use to enforce user disk space limits on the volumes used by your Windows FTP server. To set quotas, you must have Administrator rights, and the volume must be formatted with the NTFS file system. Disk quotas use NTFS' file ownership feature, which automatically attaches the name of the creator of a file to each file created on the NTFS file system.

Disk quotas are independent of the folder location of the user's files within a volume. For example, if a user moves files from one folder to another on the same volume, his space usage won't change. If he copies files to a different folder on the same volume, his space usage will double.

To set up a disk quota on, for example, a Windows 7 system, you must follow these steps:

- Open Windows Explorer. Right-click the volume to which you want to apply quotas and click Properties.
- In the Properties dialog box, click the Quota tab.
- On the Quota properties dialog box, click the Show Quota Settings button.
- In the Quota Settings for <volume_name> dialog box, check the Enable quota management box then click

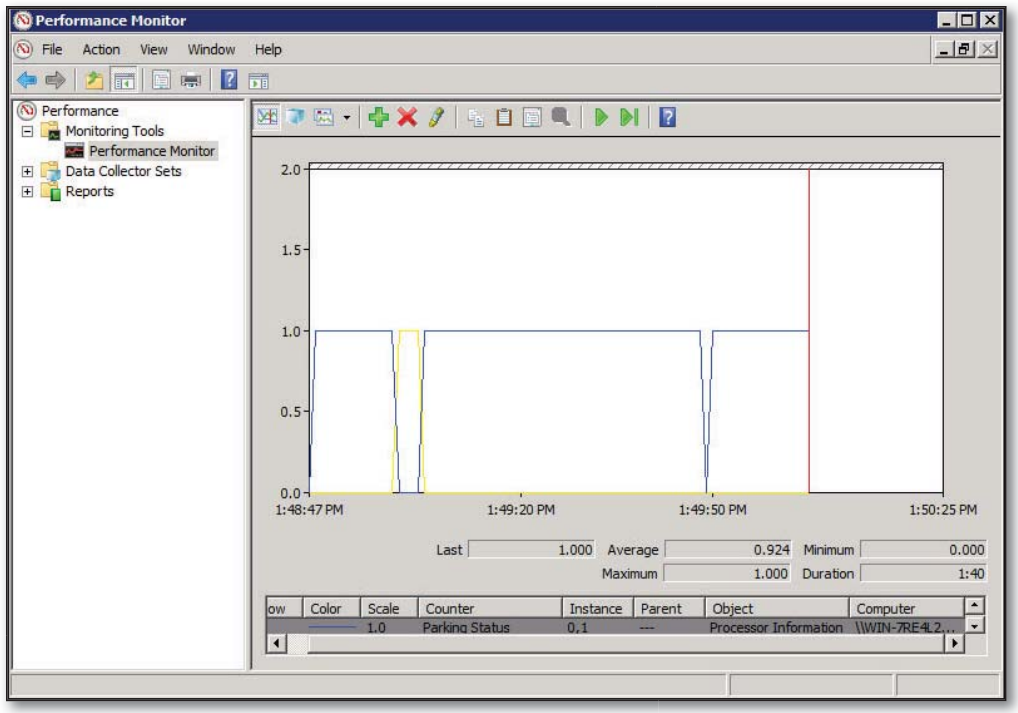

Figure 1: Performance Monitor with two cores

.exe command line tool that's bundled with Windows OSs.

—Jan De Clercq

InstantDoc ID 104688

Q: How can I view the parking state of the cores on my Windows Server 2008 R2 system?

A: Server 2008 R2 introduces the ability to put certain cores on processors to sleep if the system load doesn't warrant all the cores running. To view the core sleep state, you can use Performance Monitor and view the Processor Information, Parking Status counter for each core instance.

It's impressive how many cores are put to sleep at any one time. You can see a dual core box in Figure 1. Basically, at any time one of the cores is always asleep. I recommend changing the scale to 2 for easy viewing (viewing a value of 0 or 1 on a scale of 100 is difficult). Right-click the graph, select

Apply. If you want all new users to have access to a limited amount of disk space, click *Limit Disk Space To* then type an amount of disk space. If you want a warning message to be displayed when a user is about to reach his quota limit, click Set Warning Level To then type an amount of disk space.

- If you want to set custom disk quota for a given user, click the Quota Entries… button. This action will open a new management interface named Quota Entries for <volume_name>. From this interface, you can then define disk quota limits for specific user accounts.
- To set a disk quota limit, select New Quota Entry… from the Quota menu option and type the name of the user account for which you want to set a limit. In the Add New Quota Entry dialog box, select the Limit disk space to radio button and type a disk space limit. Also, make sure that you set a warning level.

From the "Quota Entries for <volume_name>" interface you can also easily export the quota settings for a given volume and import them into the quota settings for another volume. To do so, use the Export… and Import… options in the Quota menu.

If you want to set up disk quotas from the command line, have a look at the fsutil

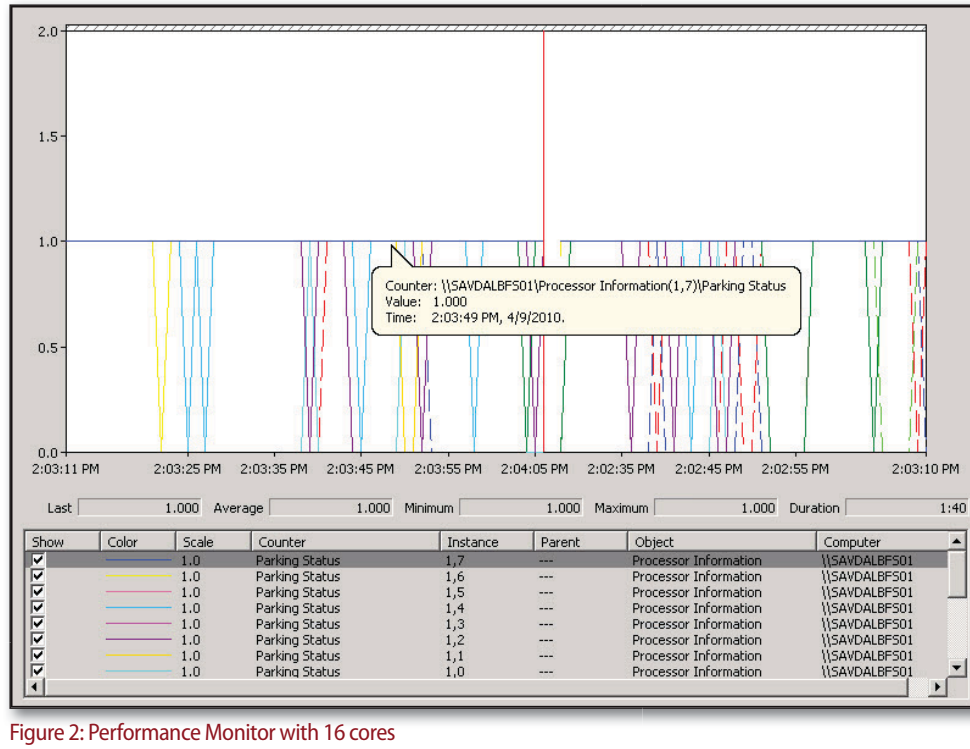

Figure 2: Performance Monitor with 16 cores

Properties, and on the Graph tab change the maximum value for the vertical scale.

On my 16-core box (eight hyperthreaded cores), which Figure 2 shows, you can see that many of the cores are asleep. By looking at the Average value for each instance, you can see what percentage of the time a core is asleep.

—John Savill

InstantDoc ID 125087

Q: Can the default encryption types the Kerberos authentication protocol uses in Windows 7 and Windows Server 2008 R2 cause compatibility problems? Is there a workaround?

A: In Windows 7 and Server 2008 R2, the DES encryption types for the Kerberos authentication protocol are disabled by default. This can cause compatibility problems if one of your legacy applications is hard-coded for only DES encryption or if the Windows account that runs a service (the service account) is configured to use only DES encryption. These services or applications will fail unless you reconfigure them to support another encryption type (RC4 or Advanced Encryption Standard, AES) or you enable DES support.

Out-of-the-box Windows 7 and Server 2008 R2 machines support the AES (to be more precise, AES128_HMAC_SHA1, AES256_HMAC_SHA1) and RC4 (RC4_HMAC_MD5) Kerberos encryption types. Microsoft only added support for the AES encryption type in Server 2008, Windows Vista, and later OSs. AES is newer and a stronger encryption algorithm than DES. The RC4 encryption algorithm has been supported by Windows Kerberos since the Windows 2000 release and is still supported in Windows 7 and Server 2008. The Kerberos logic on domain controllers will switch to AES encryption when you change your Active Directory (AD) domain to the Server 2008 domain functional level.

To check whether one of your applications or services are hard-coded to use only DES encryption, you can run a network trace when the application or service starts and check the content of the Etype fields in the Kerberos authentication headers.

To determine whether an AD user or computer account is configured for only DES encryption, you must check whether the Use Kerberos DES encryption types for this account option is set on the Account tab in the object properties (which you can access from the AD Users and Computers MMC snap-in).

If you find that you're affected by this problem, you can enable DES encryption for Kerberos authentication on Windows 7 or Server 2008 R2 using the Group Policy Object setting Network security: Configure encryption types allowed for Kerberos located in the Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options GPO container.

Microsoft has documented this problem in Knowledge Base article 977321.

—Jan De Clercq

InstantDoc ID 125072

Q: How secure is the Distributed File System-Replication (DFSR) protocol? DFSR can be used in Windows Server 2008 for replicating System Volume (Sysvol) data between Windows domain controllers. I want to know whether DFSR encrypts the data it transports across the wire.

A: Unlike other file-related transport protocols, Microsoft designed DFSR from the start with data security on the Internet and WANs in mind. DFSR communication is always encrypted. DFSR always uses authenticated and encrypted remote procedure calls (RPCs) over TCP to replicate data. It's not possible to disable DFSR's use of authenticated and encrypted RPCs.

For more information on DFSR security, see Microsoft's DFS-R specification at bit.ly/902wrN or the TechNet page "DFS Replication: Frequently Asked Questions," at bit.ly/bjsiHn.

—Jan De Clercq

InstantDoc ID 125058

Q: Why does BitLocker Drive Encryption use up most of the free space on my disk during encryption?

A: BitLocker Drive Encryption is designed to protect the enabled volumes' data, but when you delete data from a disk, you don't actually delete the content, you just remove its entries in the Master File Table (MFT). The data is still on the disk and could be read using certain utilities. It's not efficient for BitLocker to encrypt free space on a drive, so BitLocker protects this empty space by creating a large placeholder file on the drive that uses up all space except for 6GB, to keep the system running during the encryption. The data, the placeholder file, and the empty 6GB are then all encrypted. (The efficiency hit for the 6GB is accepted in the interest of security). Once the encryption is complete, the placeholder file is deleted, ensuring any remnants of data on the drive have been replaced and the free disk space returns to the correct amount.

If you need the space back, you can suspend the BitLocker Drive Encryption then resume it once you no longer need the space. (Resuming will recreate the placeholder file.)

—John Savill

InstantDoc ID 125086

Q: I made a change to my boot configuration database (BCD) or volume configuration and now BitLocker prompts for the recovery key each time I reboot. Why?

A: At least now you understand how important that recovery key is. Hope it was kept safe.

Before you make any changes to your BCD or volume configuration, always make sure you have a handy copy of your recovery key, which can be regenerated using the BitLocker Drive Encryption Control Panel applet under Manage BitLocker then Save or print recovery key again.

To stop having to use the recovery key at each reboot, you just need to suspend BitLocker then resume it. Launch the BitLocker Drive Encryption Control Panel applet, click Suspend Protection, then click Resume Protection. You'll no longer be prompted for the recovery key. You can try performing the suspend and resume before rebooting, immediately after making your BCD changes. Doing so should remove the need for the recovery key for BCD changes, but not for volume changes.

You can also suspend and resume from the command line with the commands

```
manage-bde -pause <drive>:
manage-bde -resume <drive>:
```

—John Savill

InstantDoc ID 125085

Q: Can the password of a Windows machine's domain account expire just like a normal user account's password expires (as defined in the domain password policy settings)?

A: Machine account passwords don't expire the way user account passwords do because they're exempted from the domain-level password policy and fine-grain password policies (the latter are only available in a Windows Server 2008 R2 Active Directory environment). Even if your machine has been offline for several months, it will continue to work, no matter how long it has been since its machine account password was initiated and changed.

This doesn't mean that a machine's password never changes—they're subject to another password quality control mechanism. Machine password changes are initiated from the client machine and are controlled by the local MaximumPasswordAge setting, which defaults to 30 days. When a Windows machine boots, it will notice that its password is older than 30 days and the netlogon service will initiate a password change.

If you ever encounter machine account password problems, they're typically due to the disabling or deletion of the machine account or an attempt to add a machine with the same name to the domain. In these cases, you can use the netdom.exe command line utility with the resetpwd switch to reset the machine account's password.

The netlogon registry parameters that can change the behavior of the machine password change process are MaximumPasswordAge, DisablePasswordChange, and ScavengeInterval. All three keys are located in the registry container HKLM\SYSTEM\CurrentControlSet\Services\NetLogon\Parameters.

MaximumPasswordAge determines when the password needs to be changed and defaults to 30 days. MaximumPasswordAge can be set to a value ranging from 1 to 1,000,000. In a domain, this value can be centrally controlled using the Domain member: Maximum machine account password age Group Policy Object (GPO) setting located in the Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\ GPO container.

The DisablePasswordChange key can prevent a client computer from changing its machine account password. The DisablePasswordChange key defaults to off and it's a security best practice to leave this setting untouched.

The ScavengeInterval key controls how often the netlogon scavenger thread runs. This thread is responsible for changing the machine password. ScavengeInterval defaults to 900 (15 minutes) and can be set to a value ranging from 60 to 172800 (48 hours). ScavengeInterval can also be controlled using the GPO setting Computer Configuration\Administrative Templates\System\Netlogon\Scavenge Interval.

—Jan De Clercq

InstantDoc ID 104689

Windows Server 2008 R2 and Windows 7 Group Policy

New features are more evolutionary than revolutionary

by Darren Mar-Elia

With the release of Windows Server 2008 came Group Policy preferences, a set of more than 20 Group Policy extensions that expanded the range of configurable settings within a Group Policy object (GPO). Following that game-changing release, you might expect new Group Policy features of a similar nature in Windows Server 2008 R2 and Windows 7. Unfortunately, most of what you'll see, and what I discuss in this article, are incremental improvements rather than game changers.

That being said, Microsoft did manage to incorporate one major change in Server 2008 R2 and Windows 7 by taking the first tentative steps toward automating Group Policy management using PowerShell. The rest of what you'll find new in the latest Windows release is mostly updates to existing policy areas, some additional Windows components under Group Policy management, and some improvements to Group Policy preferences. Let's look at the changes in depth.

Administrative Template Changes

The major news in Administrative Templates, or registry policy, occurred when Windows Vista shipped. With Vista, Microsoft introduced a new ADMX format and the Central Store. The ADMX format provided better multilanguage support; the Central Store took old ADM files out of the SYSVOL part of every GPO. With Server 2008 R2 and Windows 7, the greatest change in this area is the addition of yet more Administrative Template settings (more than 300). These settings cover a bevy of new Server 2008 R2 and Windows 7 features (e.g., policies to control new UI elements specific to each platform—see Figure 1). You'll find a full list of Administrative Template and Security policy settings in Excel format in Microsoft's "Group Policy Settings Reference for Windows and Windows Server" (www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=18c90c80-8b0a-4906-a4f5-ff24cc2030fb).

One of the more subtle changes to Administrative Templates is a modified ADMX schema that now supports two new registry value types: REG_MULTI_SZ and REG_QWORD. Previously, you couldn't use Administrative Templates to modify these two value types. Your choices were to deliver these kinds of values via registry scripts, or to use the Group Policy preferences' registry extension to get these value types on client machines. Now these types are supported in the ADMX syntax, and you can create custom ADMX templates that support these new types.

Another subtle Administrative Templates change is a UI improvement. In Server 2008 and Vista, Microsoft introduced the concept of Comments to Administrative Template settings. If you chose to, you could add comments to each policy setting. These comments, and the improved Explain text that provided help for each setting, were displayed as three separate tabs within Group Policy Editor's (GPE's) UI. You had to flip between each tab to use them. In Server 2008 R2 and Windows 7, all three elements are presented on a single pane that you can easily see and edit, as Figure 2 shows.

PowerShell Support

The major change in this release of Windows that I alluded to earlier is added support for PowerShell within the Group Policy universe. Microsoft added support for running PowerShell scripts within

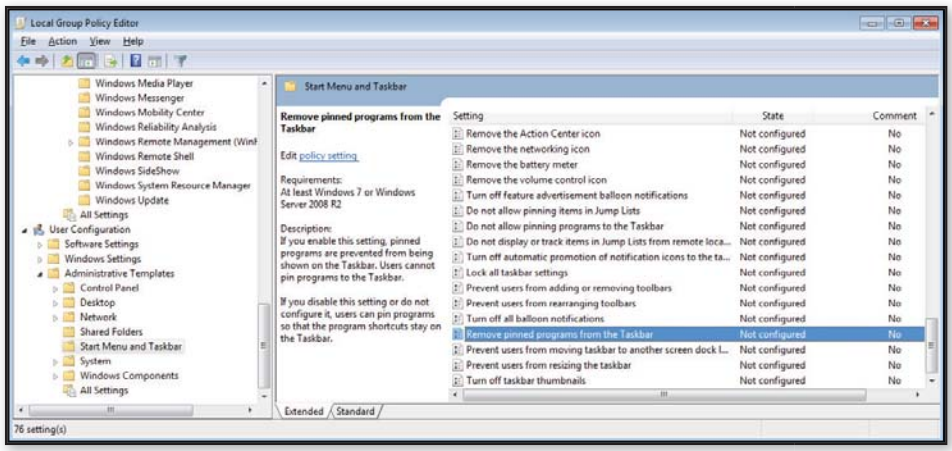

Figure 1: Server 2008 R2's and Windows 7's new Administrative Template policies

per-machine or per-user scripts policy and provided a set of 25 PowerShell cmdlets for PowerShell 2.0 that support many of the operations you can perform within Group Policy Management Console (GPMC). Let's look first at the new scripts policy support.

When you create a new startup script or logon script in GPE, you'll see a new tab. As Figure 3 shows, you can now add PowerShell scripts to your scripts policy and control whether the scripts run before or after non-PowerShell scripts. But note that only Server 2008 R2 and Windows 7 Group Policy clients will run these new PowerShell-based script policies. They won't work on earlier versions of Windows.

Perhaps the more interesting of the PowerShell enhancements is a set of cmdlets within a new PowerShell 2.0 module for Group Policy. These cmdlets encapsulate many of the functions found within the GPMC sample scripts that used to ship with that tool. From the PowerShell cmdlets, you can perform Group Policy–related administrative tasks such as creating new GPOs or deleting existing ones, linking GPOs to OUs or domains, and repermissioning GPOs. Figure 4 shows the full list of cmdlets exposed in the new module, called GroupPolicy, that ships with the Remote Server Administration Tools (RSAT) feature in Server 2008 R2 and Windows 7.

Note that to use the GroupPolicy module, you must be running PowerShell 2.0 on Server 2008 R2 or Windows 7. To provide this kind of GPMC PowerShell functionality on earlier versions of Windows, I've written a set of GPMC PowerShell 1.0 cmdlets that you can download for free at my website (www.sdmsoftware.com/freeware).

Let's look at an example of the kind of power these new cmdlets provide. Suppose you want to create, permission, and link a GPO within a PowerShell script. The following one-line command does all that by leveraging three of the new cmdlets and the PowerShell pipeline:

```
new-gpo "Marketing IT GPO" |
  Set-GPPermissions -TargetName
  "Marketing Users" -TargetType Group
  -PermissionLevel GPOEdit | new-gplink
  -order 1 -Target "OU=Marketing,
  DC=cpandl,DC=com"
```

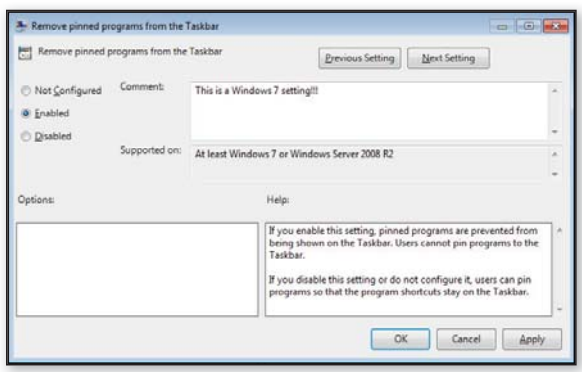

Figure 2: Unified Administrative Templates UI

In this example, I'm using three cmdlets—New-GPO, Set-GPPermissions, and New-GPLink—to create the GPO called Marketing IT GPO, to modify the GPO's permissions to grant the Marketing Users Active Directory (AD) security group rights to edit the GPO, and to link the GPO to the Marketing OU in my cpandl.com AD domain. All of this was done in one command using the PowerShell pipelining capability to pipe the output of one command to the next one. Using PowerShell and the GroupPolicy module makes it easy to perform complex Group Policy management tasks.

Modifying Registry Policy with PowerShell

GPMC functions aren't all that Microsoft put into the GroupPolicy module. Also added is basic support for modifying a small subset of the settings within a GPO. Before I describe what Microsoft has done, let me provide a little background. Currently, you have two ways in Group Policy to push out registry values. The primary way to manage registry settings in Group Policy, which we've had forever, is through the use of Administrative Templates policy. Administrative Templates use ADM or ADMX template files to build the UI in GPE that describes which registry values can be managed via Group Policy, and they also provide the help text that describes what a policy setting does. When you set Administrative Template policies in GPE, the policy settings that you define are stored within the SYSVOL portion of the GPO in a file called registry.pol. This file holds the instructions for the registry values you want to deploy in the GPO.

The other method for managing registry values is Group Policy preferences registry extension. This method is more flexible than Administrative Templates policy, and it provides the more granular targeting features that come with Group Policy preferences. Both methods will let you manage registry changes centrally, and each method has its strengths.

What Microsoft provided in Server 2008 R2 and Windows 7 with respect to PowerShell support is the ability to modify GPO settings for both of these registry methods. In the first case, you now have the Get-GPRegistryValue, Set-GPRegistryValue, and Remove-GPRegistryValue cmdlets to retrieve and modify the underlying registry.pol file that is used by Administrative Templates policy to store its settings. The advantage of using these cmdlets to modify Administrative Templates policy is that you don't need

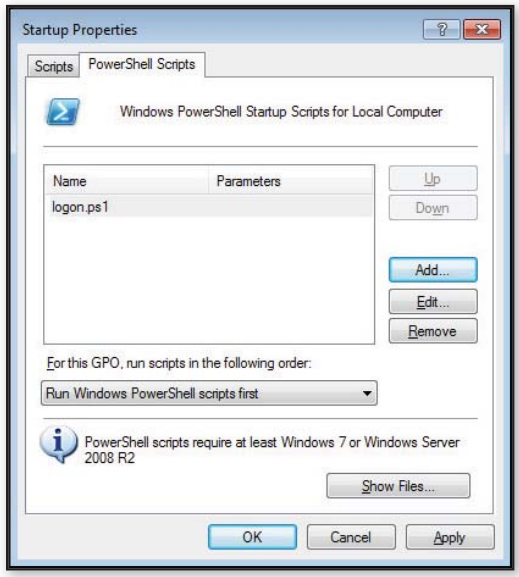

Figure 3: PowerShell support added to scripts policy

to create a custom ADM file to populate a particular registry value. Because these cmdlets don't rely on ADM or ADMX files, but instead let you push registry values into the underlying storage file, you can quickly push a new registry value into a GPO.

The downside is that this approach doesn't rely on ADM or ADMX files to define which registry values you're using; you have to know the underlying registry key and value you want to deploy. In addition, if you view a GPO that uses this method to configure a registry setting policy—and there is no underlying ADM or ADMX file representing that setting—the GPMC settings report will show it as Extra Registry Settings. If you bring up GPE, you won't see the registry value that you've pushed into that registry .pol file. Obviously, going overboard with this approach can cause confusion. I don't recommend relying on this approach for all your Administrative Template changes, but in a pinch it might come in handy.

The second PowerShell approach for registry policy involves the ability to read and modify settings in the Group Policy preferences registry extension. Microsoft ships the Get-GPPrefRegistryValue, Set-GPPrefRegistryValue, and Remove-GPPrefRegistryValue cmdlets for this purpose. These three cmdlets let you manage new Group Policy preferences registry settings (note that you can't use these cmdlets to modify existing Group Policy preferences registry policy). As an example, let's see how you can use these cmdlets to add a new registry value to a Group Policy preferences registry setting. The following example command populates the mouse beep value within a GPO called Test:

```
Set-GPPrefRegistryValue -Name Test
  -Context User -Action Create  -Key
  "HKEY_CURRENT_USER\Control Panel\
  Mouse" -Valuename "Beep" -Value "Yes"
  -Type "String"
```

Note that in this command, you have to specify the registry key, valuename, value, and type to successfully populate the Group Policy preferences setting. What these cmdlets don't provide is access to some of the more advanced features within Group Policy preferences, such as the items you see on the Common tab, and item-level targeting (the ability to granularly target preference settings). However, cmdlets are a good start toward providing Group Policy automation.

Starter GPOs—Does Anyone Care?

In Server 2008 and Vista, Microsoft introduced the notion of Starter GPOs—templates of GPOs that you could use to create real GPOs. The idea was sound, but the execution was lacking. The problem with Starter GPOs is that they support only Administrative Template policy settings, severely limiting the kinds of templates you can create. What Microsoft did in Server 2008 R2 and Windows 7 is only a small improvement over what was previously shipped. Basically, you now have the ability to prepopulate the Starter GPOs with the Windows Server 2008, Vista, and Windows XP SP2 security guideline settings that were previously provided via the so-called GPO Accelerators.

Of course, because Starter GPOs support only Administrative Templates settings, and not security settings (which are the main focus of the security guideline settings), these prepopulated Starter GPOs are relatively useless. But if you need to create templates of Administrative Templates settings that you can reuse to create real GPOs, then Starter GPOs are for you.

New Policy–Enabled Features

The last of the changes I'll cover are the new policies that have been added to support management of new features available in Server 2008 R2 and Windows 7. Most of the new policies relate to security settings, but a few minor updates have been made to Group Policy preferences as well. Let's start with the new Group Policy preferences:

- Support for managing the new Power Plans for power management that were

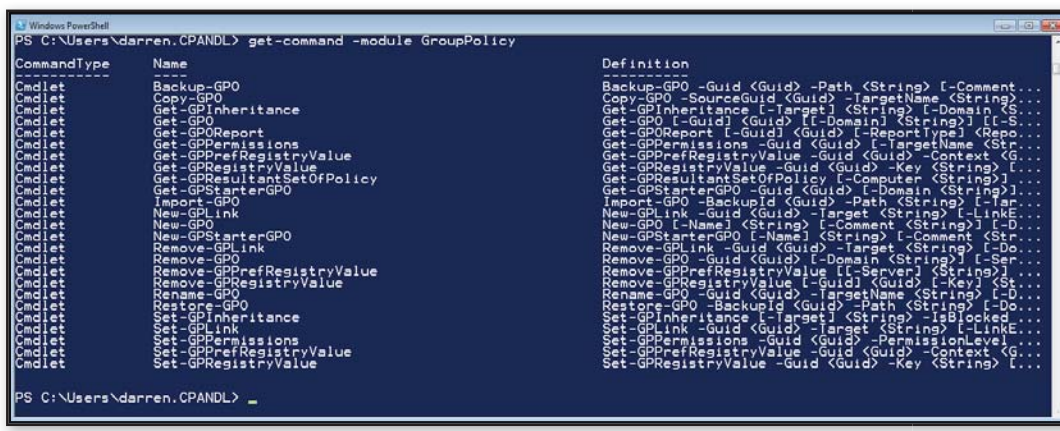

Figure 4: PowerShell cmdlets for Group Policy management

WITH MICROSOFT VIRTUALIZATION, WE

SAVED 90% IN ENERGY USAGE

BY REPLACING PHYSICAL SERVERS WITH VIRTUAL ONES

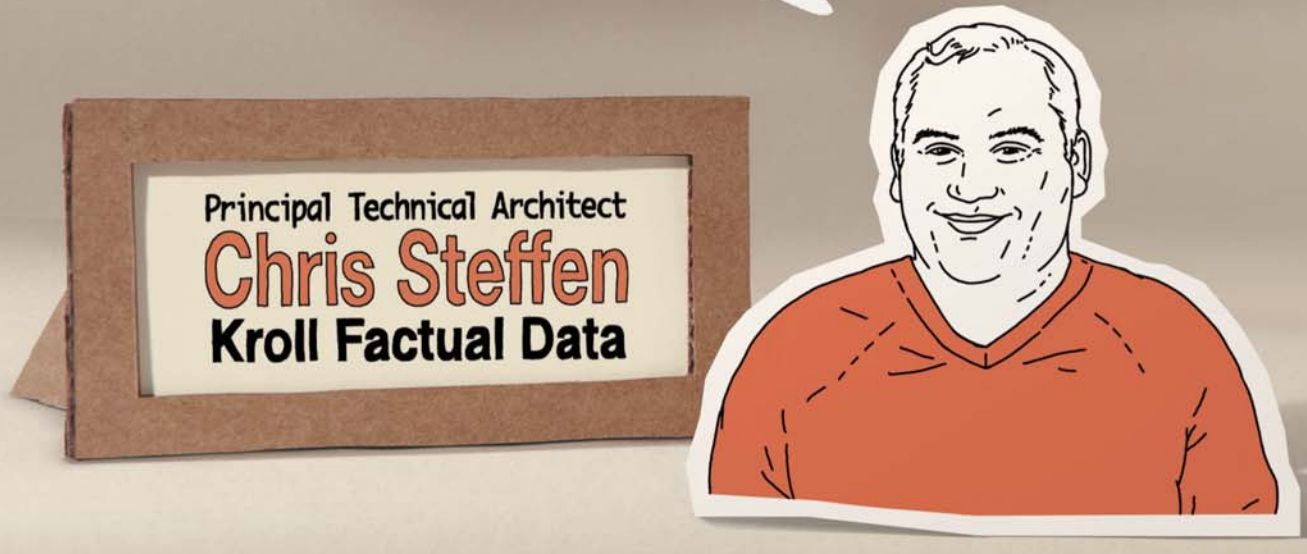

Microsoft

CASE STUDY: Kroll Factual Data

Kroll Factual Data of Loveland, Colorado, is a longtime provider of information services to the mortgage industry. The firm wanted to optimize its server infrastructure to better meet spikes in demand and reduce data center costs. Kroll Factual Data virtualized its data center using Windows Server® 2008 and Hyper-V™ technology, consolidating 650 servers to 22. It further streamlined its infrastructure using Microsoft® System Center data center solutions to monitor and manage its physical and virtual landscape, and Microsoft Visual Studio® development tools to quickly develop applications.

With its new optimized infrastructure, the company can grow faster, scale quickly to meet customer needs and dramatically reduce IT costs. Kroll Factual Data has cut annual hardware expenditures by tens of thousands of dollars, and energy costs by U.S. $442,554 annually.

To download the case study, snap this tag or text VIRTUAL to 21710*

Get the free app for your phone at http://gettag.mobi

*Standard messaging and data charges apply.

To read the full case study, visit
itseverybodysbusiness.com/virtual

introduced in Vista. These are now available in addition to Power Options and Power Schemes. Power Plans require that the client receiving them is running at least Vista.

- Updated Scheduled Tasks preferences now support the newer Task Scheduler that shipped with Server 2008 and beyond, as well as Vista. This new Task Scheduler supports many more options than Windows 2003's and XP's Task Scheduler. In addition, Microsoft added Immediate Tasks for Vista and beyond, which lets you create a one-time scheduled task that runs as soon as the policy processes.
- Addition of Internet Explorer (IE) 8 in the Internet Settings preferences, which lets you now configure options specific to IE 8.

New Security Policies

The biggest new addition in the area of Group Policy–based security policy is the Application Control Policies, or AppLocker. These policies are found under \Computer Configuration\Windows Settings\Security Settings\Application Control Policies. Essentially, this is a significant upgrade to the old Software Restriction Policies (SRPs—which are still supported in Server 2008 R2 and Windows 7) that let you control which applications can execute on your Windows systems. Specifically, AppLocker lets you create application whitelists and blacklists to explicitly allow or deny a particular application or set of applications to execute based on a set of criteria you specify.

A major difference between what's available in AppLocker and SRPs is that you now have more flexible rules for defining applications. As Figure 5 shows, for example, you can create rules by software publisher, application name, and version information held within the file.

You can also create rules for controlling script execution, which wasn't explicitly supported in earlier Windows versions. Also, for each type of rule you create, you can enforce the rule or just work in audit mode. In audit mode, whenever a rule is hit by an application, the result is logged to the client rather than blocking or allowing that application. That way, you can run a rule in test mode before making it live, to ensure it doesn't catch any unsuspecting applications. The only downside to AppLocker is that it works only on Server 2008 R2 and Windows 7 clients, so you can't leverage it in earlier versions of Windows.

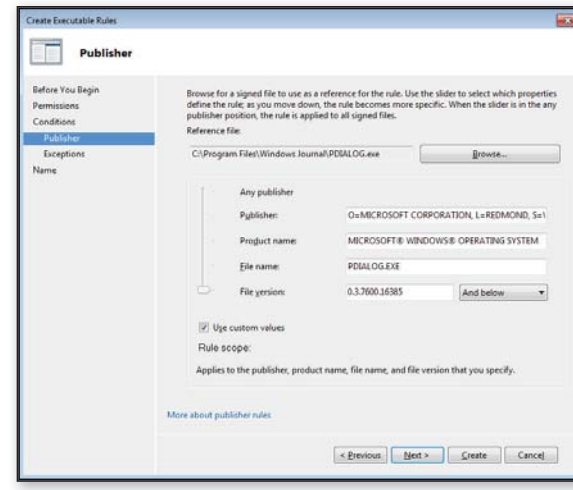

Figure 5: AppLocker rules have become more flexible

Advanced Audit Policy

Another security-related feature that you'll find in Server 2008 R2 and Windows 7 is a much more granular auditing infrastructure. If you look under \Computer Configuration\Windows Settings\Security Settings\Advanced Audit Policy Configuration, you'll see 10 different auditing categories that you can now tweak to control exactly which types of events generate security audits on Server 2008 R2 or Windows 7 systems. This new granularity, of course, is exposed only in these newest OS versions, but the fact that it's manageable via Group Policy is a good thing.

Network List Policies

The last new security policy I'll discuss gives you the ability to control network lists. By default, when Server 2008 R2, Windows 7, or Vista systems find new networks, whether public wireless networks or corporate LANs, a user is prompted to indicate the type of network it is (e.g., public, domain, home). But by using Network List Policies in Group Policy, you can now preconfigure how particular networks behave and which zone they should be shunted into when a user finds them.

You can also control the icons and the names of the networks that appear to the user. The only downside to using this policy area for preconfiguring wireless access points is that you need to know the name of the WAP ahead of time to configure all the various options. But this policy area is still a welcome addition for controlling users who frequently roam between networks.

Name Resolution Policy

The last new policy area, although not strictly a security policy (it's found under \Computer Configuration\Windows Settings\Name Resolution Policy in GPE), lets you control DNS Security Extensions (DNSSEC) and Microsoft DirectAccess DNS configurations on a per-DNS domain name basis. For example, you can configure which features of DNSSEC are used for a given client talking to its DNS server, or which DNS and proxy servers a client connecting to your network via DirectAccess will use. Although not used by all shops, this feature is handy to have in Group Policy if you're rolling out DirectAccess to your mobile users.

Evolutionary, Not Revolutionary

The Group Policy improvements in Server 2008 R2 and Windows 7 are very much evolutionary rather than revolutionary. With the possible exception of adding some PowerShell automation support for Group Policy management, this is a very ho-hum release for Group Policy fans. Sadly, we'll have to wait a while longer to see the big architectural or functional improvements that users of this more than 10-year-old technology have been looking for.

But if you're planning to roll out Windows 7, there are enough new features to make your life easier when you're configuring clients. I encourage anyone who hasn't started using PowerShell to dive in and check out what you can do with Group Policy and this new scripting technology.

InstantDoc ID 125127

Darren Mar-Elia
(darren@sdmsoftware.com) is a contributing editor for *Windows IT Pro* and is CTO and founder of SDM Software. He maintains a Group Policy resource website (www.gpoguy.com) and is coauthor of *Microsoft Windows Group Policy Guide* (Microsoft Press).

Editing and Debugging Scripts with PowerShell 2.0's Integrated Scripting Environment

The ISE is a welcome addition

by Bill Stewart

In the first version of Windows PowerShell, Microsoft didn't provide an integrated development environment (IDE) for PowerShell. In PowerShell 2.0, the PowerShell team filled that gap by adding a PowerShell IDE called the Integrated Scripting Environment (ISE). It provides an easy-to-use interface for editing and debugging scripts.

Figure 1 shows the ISE window, and Table 1 describes the elements labeled in Figure 1. The ISE window is divided into three panes by default:

- The script pane (element F) is the ISE's script editor. You can open multiple files and switch between them easily. When you open more than one file, each file gets its own tab. The script editor colorizes script code to help you identify syntax errors, automatically indents lines of code, and provides Tab-key command and path completion for paths, cmdlets, and cmdlet parameters.
- The output pane (element G) shows the output from PowerShell scripts and commands executed in the ISE.
- The command pane (element I) is an interactive PowerShell prompt.

Floating the mouse cursor over the toolbar buttons displays the buttons' functions in pop-up tooltips. The ISE's shortcut keys aren't shown in the toolbar button tooltips, but they are displayed in the menus.

All of PowerShell's Help topics are available in the ISE. Press F1 to access the topics. If you position the cursor on a cmdlet name, pressing F1 will open the PowerShell Help file to that cmdlet's Help page.

Navigating the ISE

The tabs at the top of the script pane represent open script files. When you start a new script (select File, New or press Ctrl+N) or open an existing script (select File, Open or press Ctrl+O), the ISE adds a new script tab (element D) to the top of the script pane.

In addition to script tabs, you can also open a new PowerShell tab (element C), which appears above the script tabs, as shown in Figure 1. Each PowerShell tab represents a new PowerShell instance, or execution environment, inside the ISE. This means that variables, functions, and aliases that you create in one PowerShell tab aren't visible when you switch to a different PowerShell tab. (Opening a new PowerShell tab is like starting a new powershell.exe console instance.) You open a new PowerShell tab by selecting File, New PowerShell Tab or by pressing Ctrl+T. Note that when only one PowerShell tab is open, it's not shown.

You can adjust the ISE's three panes several different ways using the View menu:

- You can select Show Script Pane Top (or press Ctrl+1) to move the script pane to the top of the ISE window. This is the default configuration.
- You can select Show Script Pane Right (or press Ctrl+2) to move the script pane to the right half of the ISE window.
- You can select Show Script Pane Maximized (or press Ctrl+3) to hide the command and output panes, maximizing the screen real estate for editing scripts.
- You can select Hide Script Pane (or press Ctrl+R) to hide the script pane, making more room for the command and output panes. Clicking the button next to element E in Figure 1 will also hide the script pane.
- You can select Command Pane Up to move the command pane above the output pane. This option doesn't have a keyboard shortcut, but you can click the button next to element H in Figure 1 to move the command pane up.

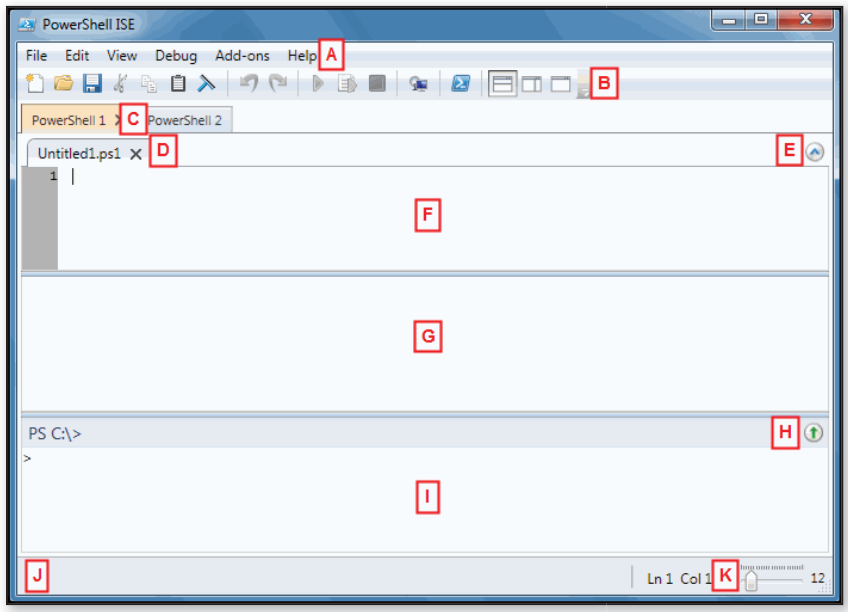

Figure 1: Identifying the UI elements in PowerShell 2.0's ISE

Opening and Editing Scripts

There are four ways you can open a PowerShell script for editing in the ISE:

1. Choose Open on the File menu (or press Ctrl+O) and select a PowerShell script from the Open dialog box.
2. Use the mouse to drag and drop a PowerShell script onto an ISE window or shortcut icon.
3. Right-click the PowerShell script in Windows Explorer and choose Edit.
4. Use the psEdit command in the command window and specify the path and filename to the script. For example, the command

```
psEdit C:\Scripts\TestScript.ps1
```

loads the file TestScript.ps1 into a new script tab. An error will appear in the output pane if the file doesn't exist.

You can open multiple scripts in the ISE using any of these four techniques. Each script will open in its own script tab. Web

Table 1: Description of the ISE's UI Elements Labeled in Figure 1

Element	Description	Keyboard Shortcut (if applicable)
A	Menu bar	Alt+first letter of menu (e.g., Alt+F if you want the File menu)
B	Toolbar*	
C	PowerShell tabs	Ctrl+Tab or Ctrl+Shift+Tab when the cursor isn't in the script pane
D	Script tabs	Ctrl+Tab or Ctrl+Shift+Tab when the cursor is in the script pane
E	Show/hide script pane button*	Ctrl+R
F	Script pane	Ctrl+I
G	Output pane	Ctrl+Shift+O
H	Move command pane above output pane button*	
I	Command pane	Ctrl+D
J	Status bar	
K	Zoom	Ctrl+Plus sign or Ctrl+Minus sign

* Floating the cursor over this item gives a pop-up description.

Listing 1: MathTest.ps1

```
# Raises a positive whole number to a
    power.
function math-power([Int] $number,
    [Int] $power) {
  if ($power -eq Ø) {
    1
  }
  else {
    $result = $number
    for ($i = 1; $i -le $power; $i++) {
      $result *= $number
    }
    $result
  }
}

math-power 2 2  # Should output 4.
```

Table 1 (www.windowsitpro.com, InstantDoc ID 104713) lists the keystroke commands available in the ISE's script editor.

Running Scripts or Selected Code

Once a PowerShell script is loaded into the script pane, you can run it by choosing Run on the File menu or by pressing the F5 key. PowerShell will run the script as if you typed the script's filename in the command pane and pressed Enter. If you want to run only a portion of a script, select the code you want to run and choose Run Selection on the File menu or press F8. PowerShell will execute the code you've selected as if you entered it in the command pane. In both cases, the output appears in the output window.

Debugging Scripts

The ISE lets you debug PowerShell scripts from its easy-to-use interface. Debugging is the process of suspending a PowerShell script as it's running so you can identify and correct problems in the code. To debug a script in the ISE, you load it into the ISE's script pane, set one or more breakpoints that suspend the script, and examine the values of variables while the script is suspended to determine the cause of the problem.

For example, MathTest.ps1 in Listing 1 is a script with a single function called math-power that contains a logic error. I'll use this script to demonstrate how to use the ISE's debugging capabilities. To follow along, you can download MathTest.ps1. Go to www.windowsitpro.com, enter 104713 in the InstantDoc ID box, click Go, click the *Download the Code Here* button, and save the 104713.zip file on your machine. After you've extracted MathTest.ps1, follow these steps:

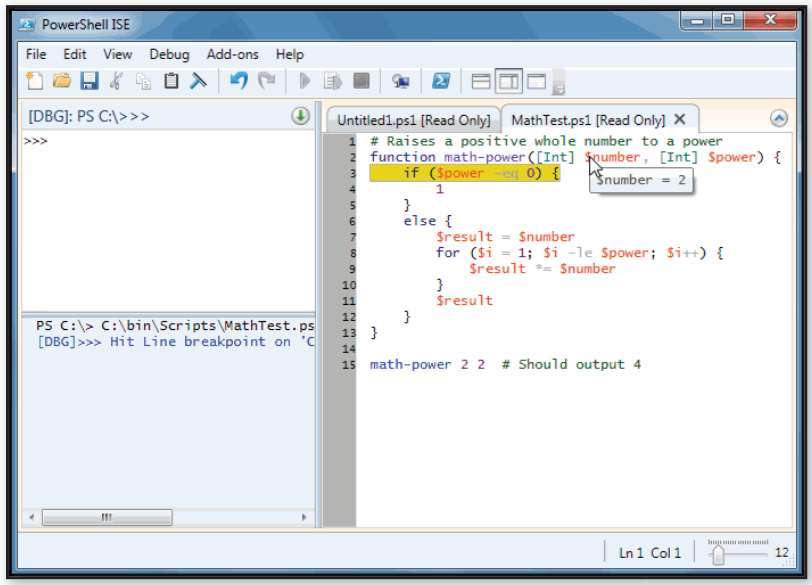

Figure 2: Displaying the variable's current value

1. Open MathTest.ps1 in the script pane by using one of the four techniques described in the "Opening and Editing Scripts" section. Press F5 to run the script. The math-power function is supposed to raise a positive whole number to a specified power. In this case, it's raising the whole number of 2 to the second power. However, notice that the result is 8 instead of 4, so there's obviously something wrong with the function.

2. Move the cursor to the first line of the function (line 3) and press F9 (or choose Toggle Breakpoint on the Debug menu). A breakpoint is a portion of code, usually a line, that suspends the script as it's running. The ISE will highlight the line with a dark red background to indicate a breakpoint is set on it.

3. Press F5 (or choose Run/Continue on the Debug menu) to run the script. The ISE will suspend the script on the breakpoint line, which will be highlighted with a dark yellow background, as shown in Figure 2. Notice how the prompt in the command pane changes to >>> and the output pane indicates that a breakpoint has been hit. While a script is suspended, you can float the mouse pointer over a variable to display the current value of that variable.

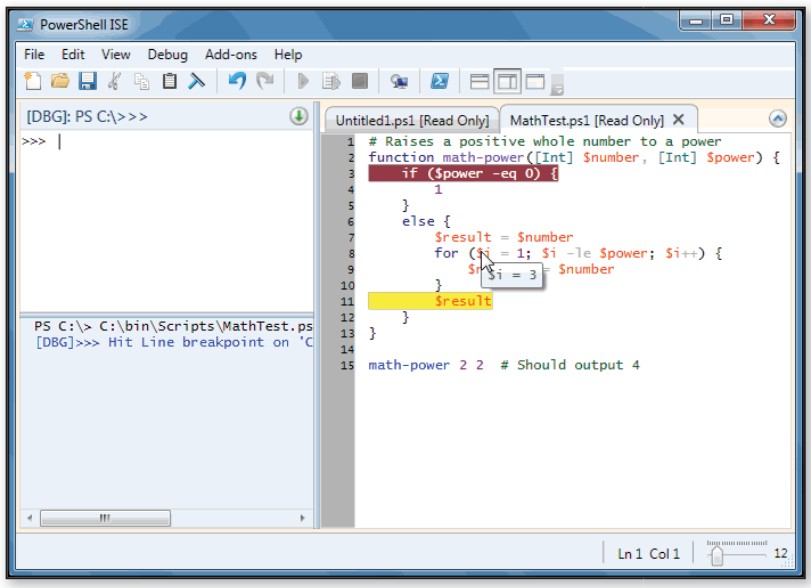

Figure 3: Finding the logic flaw

4. Press the F11 key (or choose Step Into on the Debug menu). The F11 key causes the debugger to step to the next line in the script. Because the $power variable isn't zero, the next line in the script is line 7, where the $result variable is set to $number.

5. Continue pressing F11 until the ISE pauses on line 11, where it returns the $result variable. Examine the value of the $i variable by floating the mouse pointer over it, as Figure 3 shows. You've found the logic flaw: The *for* loop has iterated three times instead of two. Press F5 to finish running the script.

6. Correct the bug by changing the initial assignment of the $i variable in the *for* loop to 2 instead of 1 and save the script. Go to line 3, press F9 to disable the breakpoint, and press F5 to run the script again. It now outputs the value 4 as it should. (The 104713.zip file includes the corrected version of the script, CorrectedMathTest.ps1.)

The debugger is much more useful with larger scripts that have many variables. However, this example should help you get started with using this helpful tool.

Add the ISE to Your Toolkit

The ISE is a truly valuable PowerShell enhancement. I recommend that you add it to your PowerShell toolkit by downloading PowerShell 2.0. This version is installed with Windows 7 and Windows Server 2008 R2, so you don't need to install PowerShell 2.0 separately if you're using those OSs. If you have Server 2008, Windows Server 2003, Windows Vista, or Windows XP SP3, you need to install the Windows Management Framework Core package from support.microsoft.com/kb/968930. Depending on your OS, you might also need to install .NET Framework 2.0 SP1 (for PowerShell 2.0) and .NET Framework 3.0 (for the ISE).

InstantDoc ID 104713

Bill Stewart
(bill.stewart@frenchmortuary.com) is a scripting guru and the lead IT analyst for French Funeral and Cremation Services in Albuquerque, NM.

Your 15 Minutes of Fame Contest

Help Windows IT Pro Celebrate 15 Years of Publication

Take 15 minutes to share a time when Windows IT Pro magazine saved your behind, made you look like a genius, or just helped you to get a good night's sleep.

Our editors and staff will select the best entry in early November. The grand prize winner will receive:

- $500 gift card
- A FREE registration to a Windows IT Pro 2011 Connections event
- A featured spot in the January 2011 issue of Windows IT Pro
- Plus, weekly drawings for $100 cash, t-shirts, and more!

Go to windowsitpro.com/go/15years now, and grab your 15 minutes of fame!

15th ANNIVERSARY Windows IT Pro 15 1995–2010

Windows IT Pro

CELEBRATING 15 YEARS IN IT WITH YOU!

Moving Mailboxes the Exchange 2010 Way

With ever-growing mailboxes in mind, Microsoft has delivered a method for online moves

by Tony Redmond

You probably know that you can't upgrade in-place to Microsoft Exchange Server 2010 from older versions of Exchange, so moving mailboxes is the only way to migrate user data to the new platform. Although the move mailboxes concept isn't new, the implementation in Exchange 2010 incorporates some innovative aspects to ease migrations. This article reviews how the Exchange 2010 move mailbox feature works and shows why the new approach taken by Microsoft makes a real difference.

Swelling Mailboxes

In all previous versions of Exchange, mailbox moves are executed within the process that initiates the move. In other words, if you initiate a mailbox move with the management console of Exchange 2007 or Exchange 2003, the move continues until the last item has been transferred from the source to the target server. During this process, the administrator can do nothing else with the console and the user is locked out of his or her mailbox because Exchange needs exclusive access to ensure that it can transfer the full mailbox content reliably. This process works and is well understood by administrators, so why change it?

The answer lies in the ever-swelling size of mailboxes. It's OK to disconnect users to move their mailboxes when the standard mailbox quota is 100MB and only a few mailboxes might grow to the giddy heights of 1GB or more. It's quite another matter when the average mailbox quota grows to 1GB and it's common to find mailboxes with quotas of 10GB or more. Indeed, in November 2009, Microsoft announced its intention to provide 25GB mailboxes to subscribers to its hosted version of Exchange ("Global Organizations Choose Microsoft Cloud Applications," www.microsoft.com/presspass/press/2009/nov09/11-02bposexpandspr.mspx) in a direct competitive move against Google, who provides 25GB mailboxes to users of Gmail for business ("Run your business, not your email server," www.google.com/apps/intl/en/business/gmail.html).

It's all very well to assign 25GB of storage to a mailbox as long as that storage never has to move—or if it can be moved without affecting the user. In fact, Microsoft designed Exchange 2010 as a platform to elegantly handle very large mailboxes, and many of the changes developers made to the Exchange Information Store are to reduce the impact on the server of multigigabyte mailboxes that contain tens of thousands of items. Changes have also been made in Microsoft Office Outlook 2007 SP2 and later

versions (including Outlook 2010, expected to be released around June 2010) to make sure that Outlook performs much better when asked to access large mailboxes.

The server and clients seem set to support large mailboxes, but as administrators are all too aware, mailbox mobility is a part of regular system management; it's a technique used to move users to new servers or to balance mailbox activity across available servers. In a world of large mailboxes, it simply becomes too slow and administratively unwieldy to disconnect users for long periods to move their mailboxes from server to server. We're moving into a world where the distinction between on-premises and cloud services blurs and companies will expect to be able to move users seamlessly between the two environments. In a nutshell, these are the reasons why the need to disconnect mailboxes before they can be moved has been replaced by a new mechanism involving move requests that are processed by the Mailbox Replication Service (MRS).

Move Requests

A move request states that an administrator wishes to move a specific mailbox from a database to another database. The databases can be on an Exchange 2010 server or on a legacy Exchange 2007 SP2 or Exchange 2003 SP2 server in the same or a different organization. Note that we don't talk about source and target *servers* anymore. The focus in Exchange 2010 is on databases because databases, not servers, are the management entity in Exchange 2010. As you probably know, Exchange 2010 databases can have multiple copies and can move between Mailbox servers. It doesn't make sense to refer to source and target servers because you don't know which server hosts the currently active copy of the database to which you want to move a mailbox.

You create move requests with Exchange Management Console (EMC) or by using the New-MoveRequest cmdlet with Exchange Management Shell (EMS). In either case, you can select a group of mailboxes and have a separate move request generated for each. You can sort mailboxes by database in EMC to make it easy to select a group from a specific database. In Figure 1, you can see I've selected a group of mailboxes to move from one database to another, then clicked New Local Move Request in the Actions pane of EMC, which invokes a wizard that generates the move requests. You can generate a group move like this if the destination database is the same for all mailboxes. If you need to move mailboxes to different databases, you have to go through the wizard separately for each destination to generate the move requests.

A *local move request* is a move to another database in the same Exchange organization. A *remote move request* is a move to a database in another Exchange organization, a feature designed to support companies that operate multiple Exchange organizations or who want to migrate from one Exchange organization to another, possibly as a result of a merger or acquisition.

The remaining steps of the wizard ask what to do if corrupt items are found in the source mailboxes. You can skip the entire mailbox and fail the move, or you can let Exchange ignore a certain number of corrupt items. The default option is to skip the mailbox, which seems extreme unless the mailbox is very new. Older versions of Exchange were less exact about item formats, and if you have mailboxes that have been used with Exchange 2003 or Exchange 2000 or have been migrated from a different email system, there's a reasonable chance that some mailboxes will contain a few corrupt items.

If you opt to skip corrupt items and Exchange encounters some during the move, they'll be ignored and you might lose some data, albeit data that can't be read by Exchange 2010. I typically set the value to 5 to let Exchange skip as many as 5 corrupt items during a move. You can check the move logs to find out if any corrupt items were found. Exchange stores details for the last two moves of a mailbox as hidden items in the mailbox's root. As you'll see later, you can access these reports with the Get-MailboxStatistics cmdlet. When everything is ready, click OK and the wizard generates the move requests for each mailbox.

Processing Move Requests

Figure 2 shows the set of move requests generated by the wizard. EMC fetches this information from Active Directory (AD) by searching for all move requests in the domain or whatever recipient scope is set for the recipient configuration section of EMC. Fetching move request data in this manner works OK for small organizations but is likely to be slow in large organizations. However, administrators working in large organizations should probably use EMS to manage move requests because the shell is faster, more flexible, and reveals more information about move requests.

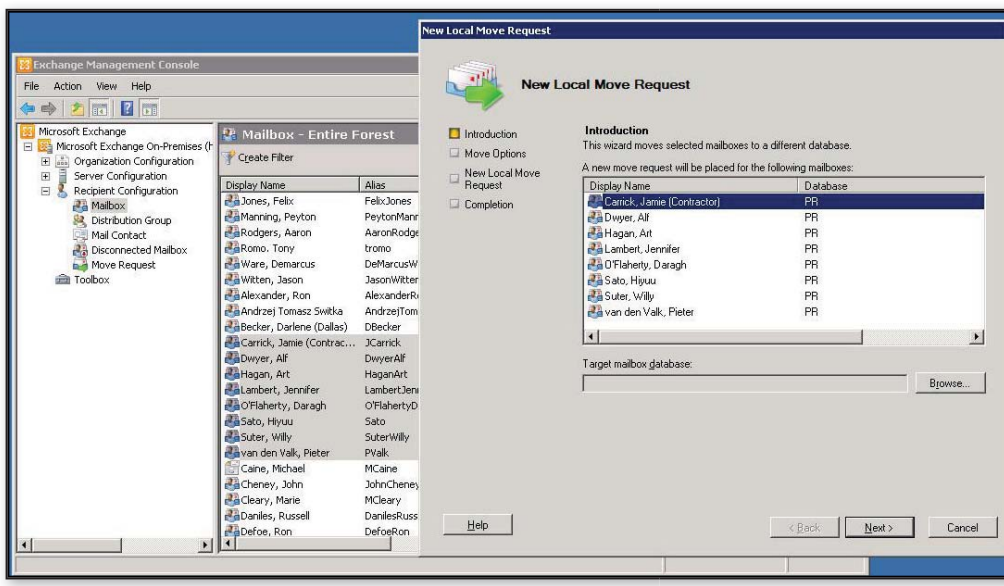

Figure 1: Creating new local move requests for a group of mailboxes

HOB RD VPN
Desktop-on-Demand

Don't Go To My PC – Go Directly To Your PC!

With HOB RD VPN Desktop-on-Demand you can access your desktop from anywhere. If your computer has been powered down, you can remotely start it.

SSL-encrypted and highly performant

The data are encrypted with SSL, and the default port 443 is used.

The RDP protocol is used for obtaining access with optimum performance.

Clientless and platform-independent No administrator rights required

This HOB software is browser-based and platform-independent, meaning you can access your data from Windows, Macs or even Linux machines.

The highly performant RDP Java client HOBLink JWT is integrated in HOB RD VPN.

Easy data transfer and local printer support

When you access your desktop, you can use the clipboard and print or transfer files over the Local Drive Mapping feature.

Desktop-on-Demand for Windows, Linux and Mac

The desktop acts as an RDP server for Windows XP, Windows Vista and Windows 7 (Exception: the Home Editions).

Even if your desktop is not running a Windows OS, HOB has a solution: HOB X11Gate for Linux or HOB MacGate for Mac OS X.

These add-on components from HOB allow you to access non-Windows desktops over the highly performant RDP protocol.

HOB RD VPN
Secure Remote Access

The Secure and Comprehensive Remote Access Software Suite!

HOB RD VPN is a software product, not a hosted service. This means your data remains fully in your hands, under your control and nobody else's.

HOB RD VPN also provides:
Remote Desktop Services (RDS)
VDI (Virtual Desktop Infrastructure)
Web Server Gate for accessing internal Web servers
File exchange with Web File Access
VT / SSH as a Java client (ideal for administrators)
HOB PPP Tunnel for universal network access
Standard emulations in Java (3270, 5250, VT, 9750)

www.hobsoft.com/DoD2

HOB RD VPN is Common Criteria certified*

* HOBLink Secure BSI-DSZ-CC-0260-2004

Figure 2: Move requests with their statuses in EMC

You'll notice in Figure 2 that the move requests have different statuses. A Queued request is one that's waiting for MRS to process it. MRS runs on every Client Access server and regularly polls AD to discover move requests in its site. By default, MRS running on a Client Access server can process up to five mailbox moves concurrently. You can tweak the MRS configuration file to force MRS to process more concurrent moves, but doing so is usually a bad idea because you run the risk of swamping target servers with the workload generated by processing incoming mailbox data. For example, Exchange updates its content indexes as mailbox data arrives on a server so that the full index is available when the mailbox move completes. It's possible there are some edge conditions that create circumstances when the MRS can process data faster, but we don't yet have enough operational experience with MRS in production environments to assess how to tweak it in different conditions.

When a move request is taken off the queue, MRS creates a copy of the mailbox in the target database and begins to transfer items. If the mailbox has an associated archive, the archive is moved too. During this time, move requests have the status Moving. If you select an individual move request and click Properties in the Actions pane, you'll see additional information, such as the size of the mailbox.

After all the data is transferred, MRS pauses to see whether any changes have occurred in the source mailbox since the move started, which is easily done by checking the last update time on the mailbox. If changes have occurred, MRS performs an incremental update to synchronize the two mailbox copies; when this process is complete, it switches the AD pointer for the user's mailbox to the newly moved copy. It's only while the pointer is switched that a user loses mailbox connectivity. Outlook Web App users get an error message stating the mailbox is temporarily unavailable, as Figure 3 shows. If they wait a few seconds and retry, Exchange will have switched over to the moved mailbox and be able to reconnect them. Desktop Outlook users receive a pop-up message to tell them that they have to exit Outlook completely and restart to connect to their moved mailbox. Microsoft would like to make this process more automatic, but the current Outlook design caches some mailbox information, such as the MAPI identifier, that has to be flushed before Outlook is able to reconnect. The situation persists in Outlook 2010.

When a move is finished, Exchange marks it with a Complete status. If you need to move the mailbox again, you have to delete the move request from EMC by selecting it and clicking Clear Move Request in the Actions pane, as Figure 4 shows. This step isn't intuitive, and Microsoft wants to incorporate some mechanism to clear up completed move requests automatically in future versions of Exchange, but for now you have to do it manually.

Scheduled Moves

EMC doesn't support scheduling move requests to occur in the future. When you create a new move request, EMC places it on the queue and MRS processes it as soon as it's ready. However, there are two ways to create a move request with EMS that allow a certain amount of control over MRS activity: You can use the -Suspend parameter or the -SuspendWhenReadyToComplete parameter. For example, you could use -Suspend as in this command:

```
New-MoveRequest
  -id 'Anil Patel' -Suspend
  -TargetDatabase 'IT Department'
```

In this instance, the queued request is on record but no data is copied to the target database until you release the move request. You would use -SuspendWhen ReadyToComplete similarly:

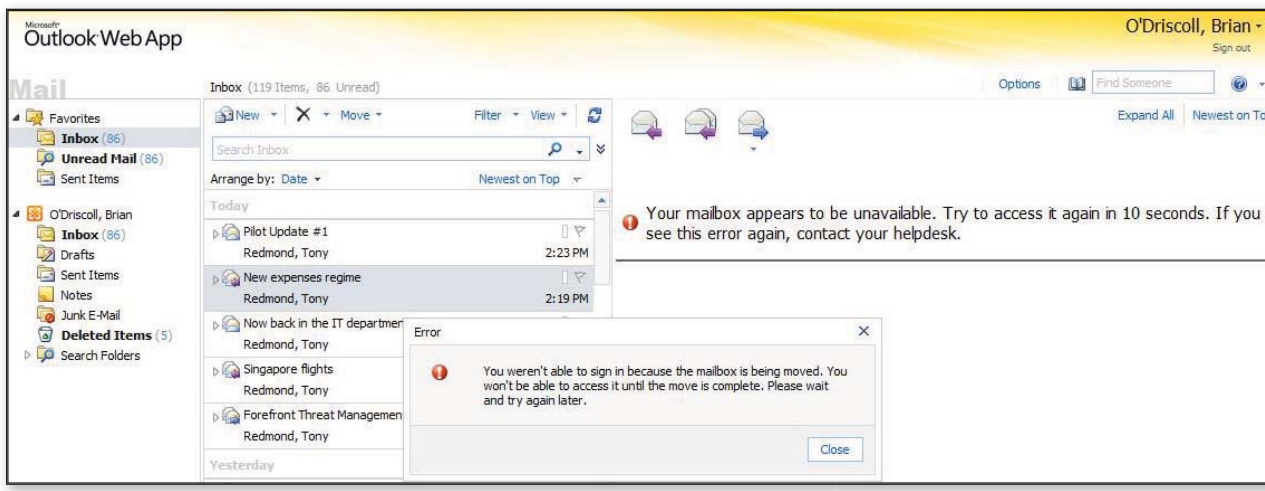

Figure 3: Outlook Web App error message during mailbox move

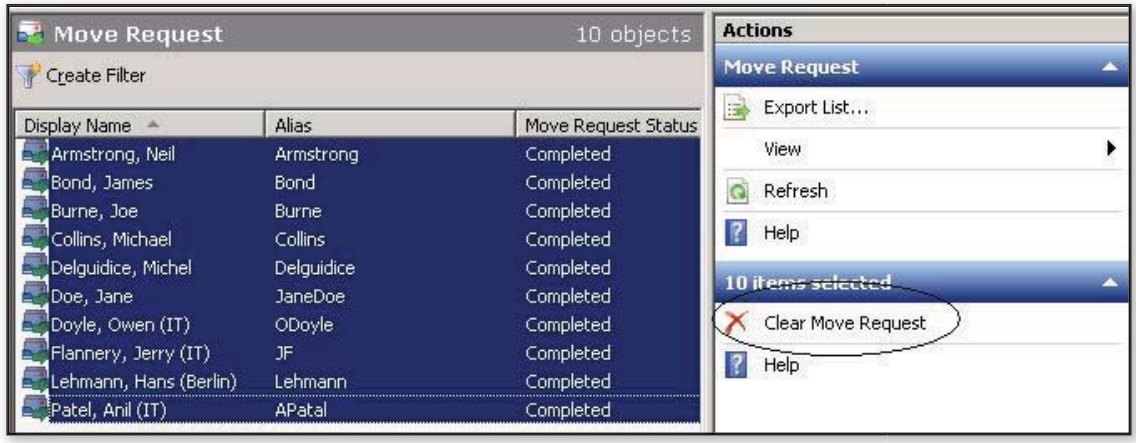

Figure 4: Clearing completed move requests in EMC

```
New-MoveRequest -id 'Anil Patel'
  -SuspendWhenReadyToComplete
  -TargetDatabase 'IT Department'
```

The big difference with this command is that it lets MRS perform the initial mailbox copy. When MRS has copied the mailbox, it puts the move request into a Suspended status and reports that a notional 90 percent of the mailbox has been copied. The last 10 percent is the incremental synchronization that MRS performs when you release the suspended move request.

You might use the first approach to create a large batch of move requests that you can release when system demand is low and the servers that support the target databases can support the additional load created by the moves. You could use the second approach to spread out mailbox moves over the working day and then perform the final switchover to the new mailbox when users are offline. The next time they connect, they access their newly moved mailbox without knowing that a move occurred.

In both cases, the move request is released with the Resume-MoveRequest cmdlet:

```
Resume-MoveRequest -id 'Anil Patel'
```

Knowing What's Happening

You can select a move request in EMC to view its properties to see what's happening as the move progresses, or you can use the Get-MoveRequestStatistics cmdlet to gain a little more insight. For example, to retrieve information about a move request, you pass the identifier for the mailbox and then select the fields that you're interested in. In the following example, I select the mailbox name, current move status, the total size of the items in the mailbox, the count of items, the percentage complete (notional) of the move, the bytes that have been transferred, and the items that have been transferred:

```
Get-MoveRequestStatistics
  -id 'Anil Patel' |
  Select DisplayName, Status,
  TotalItemSize, TotalMailboxItemCount,
  PercentComplete, BytesTransferred,
  ItemsTransferred
```

When you run this command, you might notice that there's an apparent discrepancy between the total mailbox size reported and the bytes that are actually transferred during the move. There are a couple of reasons for the difference. Along with the items that you expect to be moved, MRS moves mailbox metadata, such as the items stored in hidden folders that hold user preferences. If the mailbox is located on an Exchange 2010 server, MRS moves the contents of the dumpster (deleted items that are recoverable by users until their retention period expires). This behavior is a change from previous versions in which dumpster contents are ignored during mailbox moves. The change is necessary to support Exchange 2010's discovery search and compliance functionality. It wouldn't be very good if a mailbox move resulted in the deletion of items that might be of interest to a fraud or other type of investigation, so MRS copies the dumpster contents along with the rest of the mailbox, and content indexing captures details of these items so that they are discoverable in the new mailbox. Archive mailboxes are a new feature of Exchange 2010; if a mailbox has an archive, it gets moved too (and clearly, an archive mailbox could be much larger than the primary mailbox and so add to the time required for the overall move to complete).

Experience to date indicates that you can anticipate an Exchange 2010 server to process between 4GB and 6GB of mailbox moves per hour. The exact performance of any configuration will vary depending on the time of day, other system load, and disk performance where content indexing and transaction log generation impose a load on the storage subsystem. In any case, the performance reported by Microsoft and validated in production to date estimates that Exchange 2010 moves mailbox data at least 70 percent faster than Exchange 2007 does.

Move Reports

You can get a detailed rundown on what happens during a mailbox move by generating a complete move report with the Get-MailboxStatistics cmdlet. Here are the commands to generate a full move report:

```
$MRep = Get-MailboxStatistics
  -id TR -IncludeMoveReport
$MRep.MoveHistory[0] |
  Export-CSV 'C:\TEMP\
    MoveReport.CSV'
```

```
11/17/2009 9:00:42 AM [EX1] Connected to source mailbox 'Primary (c317cc01-ae4f-
  4bbb-aa3d-c349ad55ed5c)', Mailbox server  'ex1.xyz.com' Version 14.0 (Build 639.0).
11/17/2009 9:00:42 AM [EX1] Connected to destination mailbox 'Primary (c317cc01-ae4f-
  4bbb-aa3d-c349ad55ed5c)', Mailbox server 'ex1.xyz.com'' Version 14.0 (Build 639.0).
11/17/2009 9:00:42 AM [EX1] Move started.
11/17/2009 9:00:42 AM [EX1] Move stage: MailboxCreated. Percent complete: 5.
11/17/2009 9:00:42 AM [EX1] Move stage: CreatingFolderHierarchy. Percent complete: 10.
11/17/2009 9:00:42 AM [EX1] Source Mailbox information before the move:
Regular Items = 221 (Size: 110.5 MB (115,913,962 bytes))
Regular Deleted Items = 213 (Size: 28.33 MB (29,702,531 bytes))
FAI Items = 18 (Size: 0 B (0 bytes))
FAI Deleted Items = 0 (Size: 0 B (0 bytes))
11/17/2009 9:00:43 AM [EX1] Folder hierarchy initialized for mailbox 'Primary
  (c317cc01-ae4f-4bbb-aa3d-c349ad55ed5c)': 33 folders total.
11/17/2009 9:00:43 AM [EX1] Move stage: CreatingInitialSyncCheckpoint. Percent
  complete: 15.
11/17/2009 9:00:43 AM [EX1] Move stage: LoadingMessages. Percent complete: 20.
11/17/2009 9:00:44 AM [EX1] Soft deleted items in source mailbox:
Soft Deleted Items = 0, Size: (0 B (0 bytes))
11/17/2009 9:00:44 AM [EX1] Move stage: CopyingMessages. Percent complete: 25.
11/17/2009 9:00:44 AM [EX1] Copy progress: 0/452 messages, 0 B (0 bytes)/138.9 MB
  (145,616,493 bytes).
11/17/2009 9:00:44 AM [EX1] Messages have been enumerated successfully. 452 items
  loaded. Total size: 138.9 MB (145,616,493 bytes).
11/17/2009 9:02:30 AM [EX1] Initial seeding completed, 452 items copied, total
  size 138.9 MB (145,616,493 bytes).
11/17/2009 9:02:30 AM [EX1] Changes reported in source 'Primary (c317cc01-ae4f-4bbb-
  aa3d-c349ad55ed5c)': 0 changed folders, 0 deleted folders, 0 changed messages.
11/17/2009 9:02:30 AM [EX1] Incremental Sync 'Primary (c317cc01-ae4f-4bbb-aa3d-
  c349ad55ed5c)' completed: 0 changed items.
11/17/2009 9:02:30 AM [EX1] Move stage: IncrementalSync. Percent complete: 90.
11/17/2009 9:02:30 AM [EX1] Final sync has started.
11/17/2009 9:02:30 AM [EX1] Changes reported in source 'Primary (c317cc01-ae4f-4bbb-
  aa3d-c349ad55ed5c)': 0 changed folders, 0 deleted folders, 0 changed messages.
11/17/2009 9:02:30 AM [HPQBOX-EX2] Incremental Sync 'Primary (c317cc01-ae4f-4bbb-
  aa3d-c349ad55ed5c)' completed: 0 changed items.
11/17/2009 9:02:30 AM [HPQBOX-EX2] Move stage: FinalIncrementalSync. Percent complete: 95.
11/17/2009 9:02:31 AM [HPQBOX-EX2] Mailbox data before finalization:
```

Figure 5: Output from a Get-MailboxStatistics command

Table 1: Properties to Control MRS Operation

Property	Meaning
MaxActiveMovesPerTargetMDB	Controls the number of mailboxes that MRS can move concurrently to a single database. The default is 5 and the maximum is 100.
MaxActiveMovesPerSourceMBD	Controls the number of mailboxes that MRS can move concurrently from a source database. The default is 5 and the maximum is 100.
MaxActiveMovesPerSourceServer	Controls the maximum number of concurrent moves that a source server can perform. The default is 50 and the maximum is 1,000.
MaxActiveMovesPerTargetServer	Controls the maximum number of concurrent moves that a target server will accept. The default is 5 and the maximum is 1,000.
MaxTotalMovesPerMRS	Controls the total number of concurrent moves that a single MRS instance can process. The default is 100 and the maximum is 1,024.
FullScanMoveJobsPollingInterval	Controls the interval between MRS scans of databases looking for move requests to process. The default scan interval is 00:05:00, or 5 minutes. (MRS also processes move requests when they're initiated.)
RetryDelay	Controls how long MRS waits to retry an operation in case of transient failures. The default value is 00:00:30, or 30 seconds.
MaxRetries	Controls how many times MRS will attempt an operation in case of transient failures. The default is 60. A move fails if this value is exceeded.

You'll see that we use an object to accept piped output from Get-MailboxStatistics, then select the move history data from the object and export it to a comma-separated value (CSV) file that can be opened with Microsoft Excel or a text editor. Figure 5 shows an edited version of some of the output from this command. You can see how MRS connects to the source and target mailboxes and moves items to seed the target mailbox before pausing to check for any new items that have been created since the move began.

Managing MRS

MRS is the service that moves mailboxes from the source to target databases. Its operation is controlled by a text configuration file called MSExchangeMailboxReplicationService.exe.config, found in the Exchange binaries folder. You can edit the text file with a text editor to alter the operational parameters for MRS. Table 1 shows the MRS parameters and what they control. You have to restart the MRS process to pick up any changes that you make to these properties.

These limits are set on a per-server basis and are regarded as "best effort" limits. In other words, there might be times when the limits are exceeded. For example, two MRS servers could access the same source database to process move requests and together exceed the maximum number of active moves that you have configured the database to support. However, these instances will be exceptions rather than the rule, and normal processing will resume in due course.

Fast and Flexible Moves

Microsoft has revamped mailbox moves extensively in Exchange 2010. The new mechanism is faster and more flexible than ever before and although there are some rough edges (such as lingering completed move requests that clutter up AD), you can expect Microsoft to clean these up in future service packs.

InstantDoc ID 103651

Tony Redmond
(12knocksinna@gmail.com) is a contributing editor for *Windows IT Pro* and author of *Microsoft Exchange Server 2007 with SP1* (Digital Press).

Managing Privileged Access to Servers

Techniques for managing administrator access to Windows servers

by Russell Smith

Even for small teams working with a handful of servers, managing privileged access can make the difference between stable, secure systems and uncontrolled change that imperils a company's systems and data. As networks grow, the need to manage privileged access to servers as a means of basic security and change control is simply unavoidable and might also be a prerequisite for regulatory compliance. Let's dive into the access problems that many companies face, then walk through some basic steps that can put your organization on the right path to more secure systems.

Defaulting to Domain Admins

Small enterprises frequently hand out accounts with Domain Admin privileges to IT staff. Not only is this the easiest way to give immediate root access to all workstations and servers (because the Domain Admins group is added to the local Administrators group when computers are joined to a domain), but it also provides write/change access to objects stored in Active Directory (AD). But how often do junior administrators or support professionals really need Domain Admin privileges? Apart from the risk involved in granting new staff members root privilege to your entire network on their first day, administrative access to workstations needn't be via the Domain Admins group. Best practice is to provide root access to servers on an as-needed basis.

Determining Access Levels

In contrast with UNIX security, which uses ACLs to control access to files, Windows secures all objects and comes with a series of built-in groups intended to ease the permission-granting process. Although it might be tempting to take advantage of the more granular security model in Windows and add sysadmin accounts to these groups, it should be done with caution as it becomes more difficult to audit system changes. However, it might be permissible to make a sysadmin a permanent member of read-only groups, such as Event Log Readers, as those groups don't grant any administrative access to servers.

Service accounts often require administrative access or a privilege that standard users don't hold. Such accounts can also be used to run scheduled tasks or batch files and might also require elevated privileges. Accounts used to give temporary access to servers (i.e., firecall accounts) for maintenance purposes also require privileged access. Any other requests for permanent privileged access to servers outside the scope of service and firecall accounts should be carefully considered, and in many cases is unlikely to be necessary.

Setting the Scene

Before we look at working directly with servers and privileged accounts, it's important that we have an AD structure that facilitates easy server management. The organizational unit (OU) hierarchy that I'll illustrate in this article is based on recommendations in the *Windows Server 2008 Security Compliance Management Toolkit* (technet.microsoft.com/en-us/library/cc514539.aspx) and represents a typical structure deployed in organizations. In Figure 1, the GPOs with the prefix *WS08 EC* were automatically generated using the toolkit's GPOAccelerator tool and contain the recommended settings for each server role in a standard enterprise configuration.

You'll notice that in addition to the default container for domain controllers (DCs), I've created OUs for Firecall Accounts, Service Accounts, and WS08 EC Member Servers. Using OUs to separate firecall and service accounts from standard domain user accounts lets you delegate control of specific account properties in the OU to a limited set of users.

The WS08 EC Member Servers OU is further divided into OUs for specific server roles: WS08 App Servers and WS08 EC File Servers. The WS08 App Server Baseline GPO was created using the Security Configuration Wizard in Windows Server 2008 as these servers are specific to my organization. The WS08 App Servers OU contains the DMZ OU, so different security settings can be applied according to the servers' exposure to external threats.

Firecall Accounts

Before disabling the built-in administrator accounts and randomizing the password, you should create a domain user account for each server in a dedicated Firecall Accounts OU. We'll use the Restricted Groups policy to add the appropriate firecall account to each server. Firecall accounts are managed by a dedicated security team, who enable the otherwise disabled accounts and issue randomly generated passwords to sysadmins for a limited time, after which accounts are disabled and passwords changed.

Every time a firecall account is enabled and a password is issued to a sysadmin, the security team manually logs the details so that any problems resulting from changes made during the logon session can be traced back to a specific user, time, and change request. If sysadmins use their own accounts to log on to servers, we can track user logon and logoff events but not the level of privilege. Ideally, the security team wouldn't know the passwords issued for firecall accounts, but that would require the deployment of a dedicated password-management system.

Delegating Control to Manage Firecall Accounts

In my example, I've created four firecall accounts (e.g., AppServ01_Firecall, AppServ02_Firecall) for my application servers in the Firecall Accounts OU. I have a specific team of security administrators whom I'll allow to enable, disable, unlock, and reset the passwords on accounts in the Firecall Accounts OU.

Let's use the Microsoft Management Console (MMC) Active Directory Users and Computers snap-in to delegate specific access to the Firecall Accounts OU:

1. Under Administrative Tools, open Active Directory Users and Computers as a domain administrator. In the left pane, expand your domain, right-click the Firecall Accounts OU, and select Delegate Control from the menu.
2. In the Delegation of Control Wizard, click Next on the welcome screen.
3. On the *Users and Groups* screen, click Add. In this example, I have a predefined group called On Duty Security Administrators, to which I want to delegate control of user accounts in the Firecall Accounts OU. You can use any domain group you choose. In the *Select Users, Computers, or Groups* dialog box, type the name of the group you want to delegate control to, and click OK. Click Next on the *Users and Groups* screen.
4. On the *Tasks to Delegate* screen, select *Create a custom task to delegate* and click Next.
5. On the Active Directory Object Type screen, select *Only the following objects in this folder* and select *User objects* at the bottom of the list. Click Next.
6. On the Permissions screen, select General and Property-specific. In the list, select the following permissions: *Change password, Reset password, Read lockoutTime, Write lockoutTime, Read userAccountControl*, and *Write userAccountControl*. Click Next, then Finish.

Any users in the On Duty Security Administrators group who are otherwise standard domain users will now be able to open Active Directory Users and Computers and perform limited actions on user accounts in the Firecall Accounts OU.

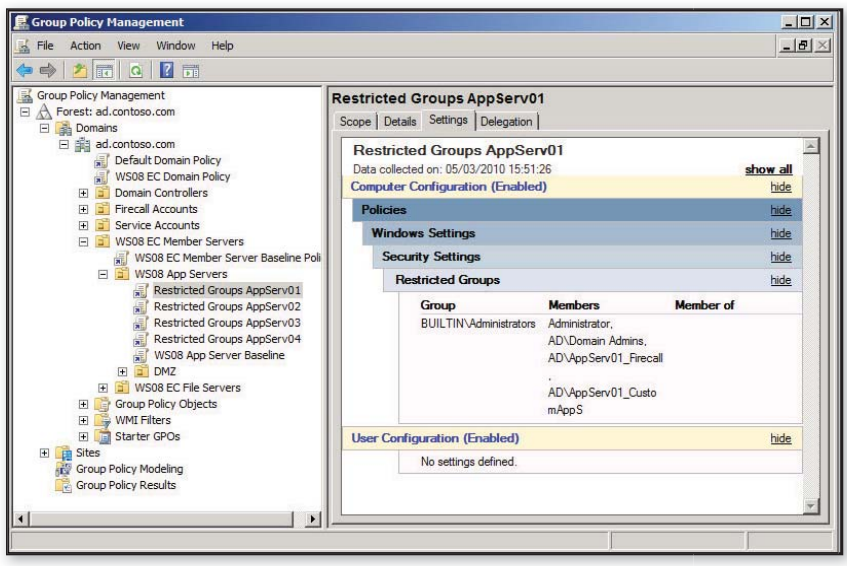

Figure 1: Optimal OU hierarchy

Managing Privileged Server Access with Restricted Groups Policy

In Figure 1, apart from the WS08 App Servers Baseline GPO, there are four additional GPOs linked to the WS08 App Servers OU prefixed with Restricted Groups that are specific to the four servers whose computer accounts are located in the OU. Each Restricted Groups GPO contains a WMI filter to ensure that the settings apply only to the relevant server. To create a Restricted

Groups GPO, open Group Policy Management from Administrative Tools on the Start menu as a domain administrator:

1. In the left pane of Group Policy Management, expand your forest and domain, as the figure shows.
2. Right-click the Group Policy Objects container and select New from the menu.
3. In the New GPO dialog box, name the GPO Restricted Groups, add the name of a specific server (e.g., *Restricted Groups MyServer1*), and click OK.
4. Expand the Group Policy Objects container, right-click the new GPO in the list, and select Edit from the menu.
5. In the Group Policy Management Editor, expand Computer Configuration, Policies, Windows Settings, and Security Settings.
6. Right-click Restricted Groups, and select Add Group from the menu.
7. In the Add Group dialog box, type *Administrators* in the Group field and click OK.
8. In the Administrators Properties dialog box, click Add to the right of *Members of this group*.
9. In the Add Member dialog box, type *Administrator* in the *Members of this group* field.
10. Click Browse, and use the Select Users, Service Accounts, Groups dialog box to add the Domain Admins group and optionally any service or firecall accounts specific to the server. It's important to note that the Restricted Groups policy doesn't amend local group membership, but it replaces the existing members each time Group Policy is refreshed on the server. Therefore, you need to include groups and accounts that already exist on the server. Restricted Groups policy shouldn't be used on DCs to manage domain-based groups.
11. Click OK when you've located all the domain accounts and/or groups that you want to add. In the Add Member dialog box, the selected accounts will appear in a semicolon-delimited list. Click OK to continue.
12. Click OK in the Administrators Properties dialog box, and close the Group Policy Management Editor.

Before linking our new GPO to an OU, we need to add a WMI filter to ensure that the policy applies only to a specific server. Alternatively, you could create a separate OU for each server. Either way, an additional AD object is necessary.

1. Back in Group Policy Management, right-click WMI Filters in the left pane and select New from the menu.
2. In the New WMI Filter dialog box, call the filter *MyServer1*, add a description, and click Add to the right of Queries.
3. In the WMI Query dialog box, type *Select Win32_ComputerSystem where Name = "MyServer1"* in the Query field and click OK.
4. Back in the New WMI Filter dialog box, which Figure 2 shows, click Save.
5. Under Group Policy Objects, select the GPO we created earlier.
6. In the right pane, switch to the Scope tab. In the WMI Filter section, use the drop-down menu to select the just-created MyServer1 filter. Click Yes in the warning dialog box to apply the filter.

You can now link the GPO to the appropriate OU to apply the Restricted Groups policy to the server. If you try to add an account to the local Administrators group on the given server, then refresh the policy by using the Gpupdate /force command, you'll notice that the account you added is removed.

Built-in Administrator Accounts

The administrator account should be disabled and password randomized when machines are initially deployed. Although the built-in administrator account is disabled by default in Server 2008 and Windows Vista or later, the password for this account should be randomized on each server as it's still possible to use this account to boot into Safe Mode and log on. For already-deployed devices, you can use a script to disable the built-in administrator account, as well as change and record the password. The BIOS of each server should be protected by a password and the boot order set to start only from the local disk.

Service Accounts

Service accounts often require privileged access, and like firecall accounts, they're specific to a given server. As such, service accounts can be managed simlarly to firecall accounts in their own dedicated OU with control delegated to a limited set of people.

Server 2008 R2 introduced the concept of Managed Service Accounts (MSAs) and Virtual Accounts. MSAs are domain accounts created and managed using PowerShell cmdlets that have automatically generated 240-character passwords that are reset every 30 days by the Netlogon service in much the same way as computer accounts. MSAs can be associated with only one computer (Server 2008 R2 or Windows 7), but multiple MSAs can be used on a single machine. If your AD forest is running at Server 2008 R2 functional level, you can see the Managed Service Accounts container in Active Directory Users and Computers if you select Advanced Features on the View menu.

Virtual Accounts aren't created but are used to run services. You access them by typing *NT SERVICE\[service name]* into the *This account* field on the service's Log On

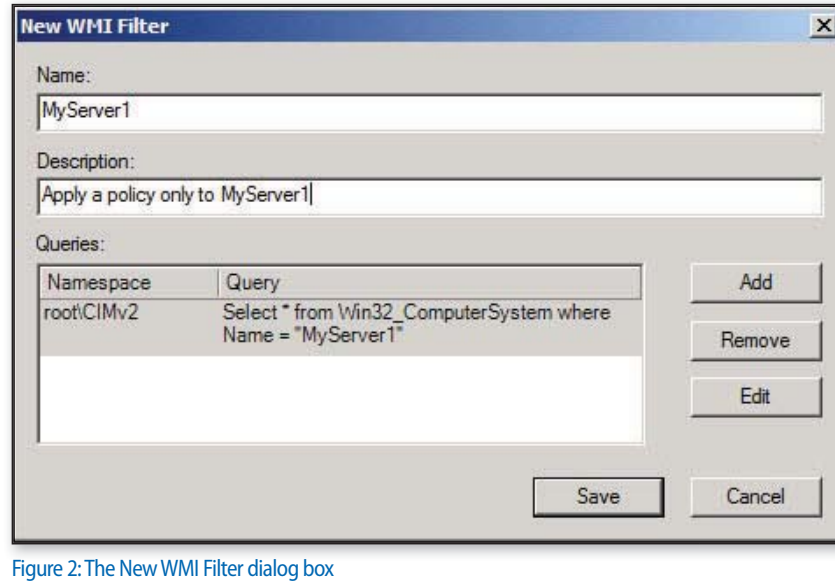

Figure 2: The New WMI Filter dialog box

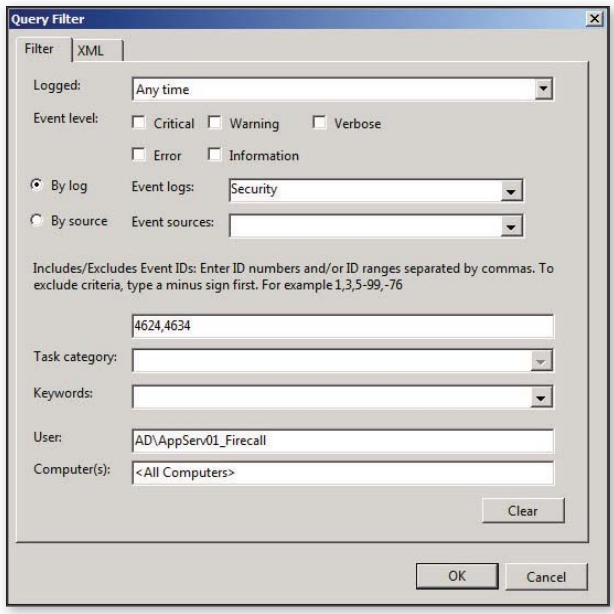

Figure 3: Forwarding only the relevant logon/logoff events for a given user

tab in the service management console and leaving the password fields blank. Virtual Accounts have the same permissions as Network Service: They have standard user rights on the local machine but communicate with the network as the local computer account. Virtual Accounts should be used if no specific access to the domain is required. For more information about MSAs and Virtual Accounts, see "Use MSAs to Ease the Pain of Administering Service Accounts" (InstantDoc ID 103625).

Auditing Logoff and Logon Events

Once a sysadmin informs the security team that his or her work is complete on a server and the firecall account can be disabled, it might be worth checking the server's security event log to ensure that the firecall account has been logged off. Even if an account is disabled, that doesn't necessarily stop the account from being used to perform administrative functions on a system if it has already received an access token. You can give the security team read access to security event logs on servers and check for 4624 and 4634 logon/logoff events.

You can also use event forwarding and collect logs on a central server. If the source servers are running Server 2008, the Network Service account must be added to the local Event Log Readers group so that security events can be forwarded. For more information about the requirements for forwarding security events, see "Forwarding Security Events from Windows XP, Server 2003, and Vista/Server 2008" (blogs.technet.com/wincat/archive/2009/06/23/forwarding-security-events-from-windows-xp-server-2003-and-vista-server-2008.aspx).

To read the Forwarded Events log on the server collecting events, add the appropriate users or groups to the server's Event Log Readers group. Each subscription can be filtered so that only the relevant logon/logoff events for a given user are forwarded, as you see in Figure 3.

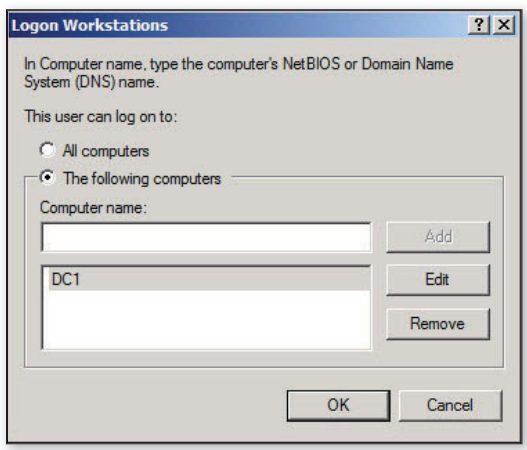

Figure 4: The *This user can log on to* feature

Despite the fact that Restricted Groups policies are a fairly robust way to ensure that only specified accounts and groups are members of the local Administrators group on servers, you could also use a script to audit local Administrators and compare the results with a known membership list for each server. The following PowerShell code uses ADSI to enumerate members of the local Administrators group on a machine called MyServer1:

```
$group =[ADSI]"WinNT://
   MyServer1/Administrators"
$members = @($group.psbase
   .Invoke("Members"))
$members | foreach {$_.GetType()
   .InvokeMember("AdsPath","GetProperty",
   $null,$_,$null)}
```

Domain Controllers

Privileged access to DCs presents a particular problem because there are no local groups, and administrative access granted to a user via a built-in AD group (e.g., Server Operators) ultimately provides enough access for a user to elevate privileges, making it impossible to completely separate access to the server and AD.

Nevertheless, firecall accounts can be created for DCs. Due to the distributed nature of AD's built-in groups, these accounts could be restricted to the DC, which they can log on to by using the *This user can log on to* feature on the Account tab of the user's properties in Active Directory Users and Computers, as you see in Figure 4. Ultimately, once a user has Domain Admin privileges, any restrictions imposed can be circumvented. Everyday administrative tasks that require a connection to a DC should be carried out using the Remote Server Administration Tools (RSAT) and the appropriate permissions delegated to users.

Minimizing Risk

Establishing simple procedures and delegation of control can significantly decrease the risk associated with IT staff maliciously or unwittingly destabilizing systems or gaining unauthorized access to company data. Although the techniques in this article might not be enough to satisfy auditors, they should put you on the right track. For larger organizations, third-party products such as BeyondTrust's PowerKeeper and ManageEngine's Password Manager Pro help fully automate management of privileged accounts across hundreds of servers and assist in achieving compliance with government or industry regulations.

InstantDoc ID 104709

Russell Smith
(rms45@rsitc.com) is an independent IT consultant. He has been working in IT since 2000, specializing in systems management and security.

Exchange Server's Client Access: An Introduction

Learn how this essential role handles all your mailbox connections and how to plan and size your Client Access infrastructure

by Ken St. Cyr

The Client Access server role, which debuted in Microsoft Exchange Server 2007, gets a significant increase in responsibility in Exchange Server 2010: It now handles mailbox access for all client connections. Because of the emphasis that's being placed on the Client Access role, it makes sense to gain a better understanding of the role and how to plan for its deployment. This article is the first of a series that will take an in-depth look at this vital Exchange role. I'll start by helping you understand what the Client Access role does and what you should consider when planning your Client Access server infrastructure. In future articles, I'll walk you through the deployment and configuration of your Client Access servers, show you how to make them resilient against failures, make them more secure, and help you keep them healthy.

Overview of the Client Access Role

The Client Access server role is one of five potential server roles in Exchange. It can be combined on the same server with the Mailbox, Hub Transport, and Unified Messaging roles. However, regardless of the combination of roles, you'll need to ensure that the hardware (processor, memory, hard disks, and network connections) is sized to handle the workloads of each role. I'll explore sizing considerations later in this article.

Also, depending on which roles you group together, you might be limited in which security policies and hardening measures you can apply to the server. For example, if you have the Mailbox role and Client Access role combined on the same server, you make the Mailbox role vulnerable because it will have additional services installed and exposed. Each role has its own set of responsibilities, and the Client Access server doesn't exactly have a light load. Many people think of the Client Access role as only providing the interface for web-based email, but in reality it does much more, particularly in Exchange 2010. But before we dive into Exchange 2010 specifically, I'd like to walk you through the basic functionality that the Client Access server brings to the table.

In Exchange 2007, the Client Access server role was introduced to remove some of the burden from Mailbox servers and to provide a common entry point into the organization for client protocols, including connections from web browsers, mobile phones, or even email clients over the Internet. As a result, resources on the Mailbox server role are freed up so they can be focused on working with the messaging data. Also, your Mailbox servers are more secure because it's less likely that administrators will run services on them that are used by external clients—but hopefully you were doing this part anyway with a reverse proxy server.

To provide these connection points, the Client Access server uses a mixture of Windows services and web virtual directories provided through Microsoft Internet Information Server (IIS). You can see the virtual directories provided by the Client Access server by taking a look in the Connections (left) pane of the IIS Manager tool, which Figure 1 shows. Some of these services probably look familiar, but you might wonder what the others do. Table 1 summarizes the virtual directories that the Exchange 2010 Client Access server runs.

Client Access in Exchange 2010

The Exchange 2010 Client Access server brings with it a variety of new features, some improvements to existing features, and some well-needed interface enhancements to make the Exchange administrator's job easier. There are two new virtual directories in the Exchange 2010 Client Access server that build on the existing web-based architecture used in Exchange 2007. First, the Exchange Control Panel (ECP) provides a new interface for browser-based email clients. ECP works in conjunction with Outlook Web App (OWA; formerly Outlook Web Access) to give users more control over their mailbox settings.

Users can change their contact information in the Global Address List (GAL) and manage distribution groups they own through ECP. If users have administrative roles assigned through Exchange's implementation of Role Based Access Control (RBAC), they can perform additional management tasks for the Exchange organization, such as performing discovery searches across multiple mailboxes and assigning roles to other users. As Figure 2 shows, these additional capabilities show up as added tabs in the ECP UI.

The second new virtual directory introduced in the Exchange 2010 Client Access server is Remote PowerShell, which lets you connect to Exchange Management Shell (EMS) from any computer with an SSL connection to the Client Access server. Because this communication happens over HTTPS, you can even use computers outside your network. Although you're running commands from your client, they're actually executed from the Client Access server. Therefore, you can use Remote PowerShell to manage your 64-bit Exchange servers from your 32-bit Windows 7, Windows Vista SP1, and Windows XP SP3 clients.

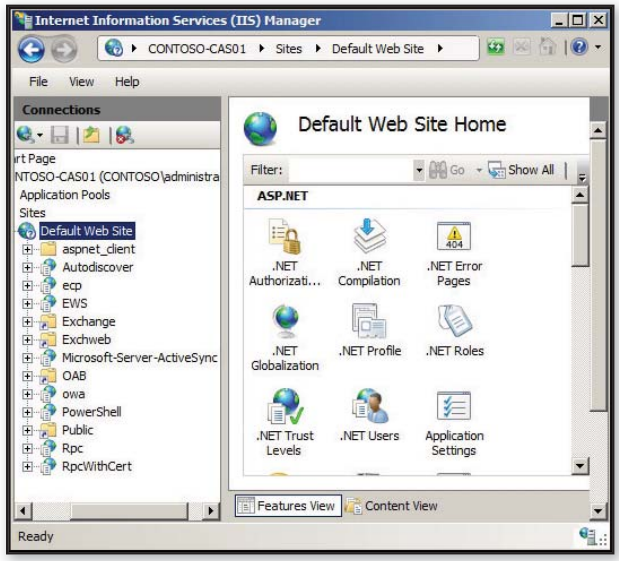

Figure 1: Client Access server virtual directories revealed in the Connections pane of IIS Manager

The prerequisites for your clients to access Remote PowerShell are for Windows Remote Management (WinRM) 2.0 and Windows PowerShell 2.0 to be installed. These components are provided in the Windows Management Framework (support.microsoft.com/default.aspx/kb/968929). After you install the Windows Management Framework on your clients, you don't need to install the Exchange Management tools on them, nor do the clients need to be joined to the domain. You authenticate through Remote PowerShell by creating a session object and specifying the PowerShell web services Uniform Resource Identifier (URI), as the following code shows:

```
$Session = New-PSSession
  -ConfiguratonName Microsoft.Exchange
  -ConnectionUri https://contoso-cas01
  .contoso.com/PowerShell/
  -Authentication
    NegotiateWithImplicitCredential
Import-PSSession $Session
```

Although the code appears on multiple lines here due to space constraints, you would enter it all on one line. Along with Kerberos authentication (which requires a domain connection), Remote PowerShell supports NTLM, basic and digest authentication, and the Credential Security Support Provider (CredSSP) protocol.

The Client Access server determines which RBAC roles you hold and presents

Table 1: Virtual Directories on Exchange 2010 Client Access Servers

Virtual Directory	Description
Outlook Web App (OWA)	OWA is the web-based email client.
Availability	The Availability service distributes free/busy information about users in the organization.
Offline Address Book (OAB)	The OAB virtual directory provides the OAB for your Outlook clients to download. In Exchange 2003 and earlier, both the OAB and free/busy information were distributed solely through system public folders. The Client Access server added web-based distribution. Outlook 2003 clients don't support web-based OAB distribution, so they still require public folder distribution.
Exchange Web Services (EWS)	EWS provides a web services API for your applications to access Exchange data.
Outlook Anywhere	Outlook Anywhere gives your users the ability to connect to their mailbox with Outlook from the Internet. Outlook clients use remote procedure calls (RPCs) on your intranet to connect to Exchange. To use RPC without Outlook Anywhere, you need to open a wide range of ports in your firewall because RPC dynamically allocates ports for the connection. However, Outlook Anywhere gives external Outlook clients access by wrapping the RPC communications in an HTTPS connection, which requires only port 443 to be open.
Autodiscover	Autodiscover lets your Outlook clients perform automatic configuration. Rather than having your users manually type in their Exchange server settings, Autodiscover determines what this information is so the user doesn't need to know it.
Exchange ActiveSync (EAS)	EAS provides connectivity for ActiveSync-enabled mobile devices.
Remote PowerShell	Remote PowerShell lets administrators remotely execute PowerShell commands from their WinRM 2.0 and PowerShell 2.0 clients. This service is new to Exchange 2010.
Exchange Control Panel (ECP)	ECP is a new component in Exchange 2010. It replaces the Options page in OWA and gives end users some additional options for managing their contact information and groups. For administrators, ECP provides an additional interface for managing some organizational settings.

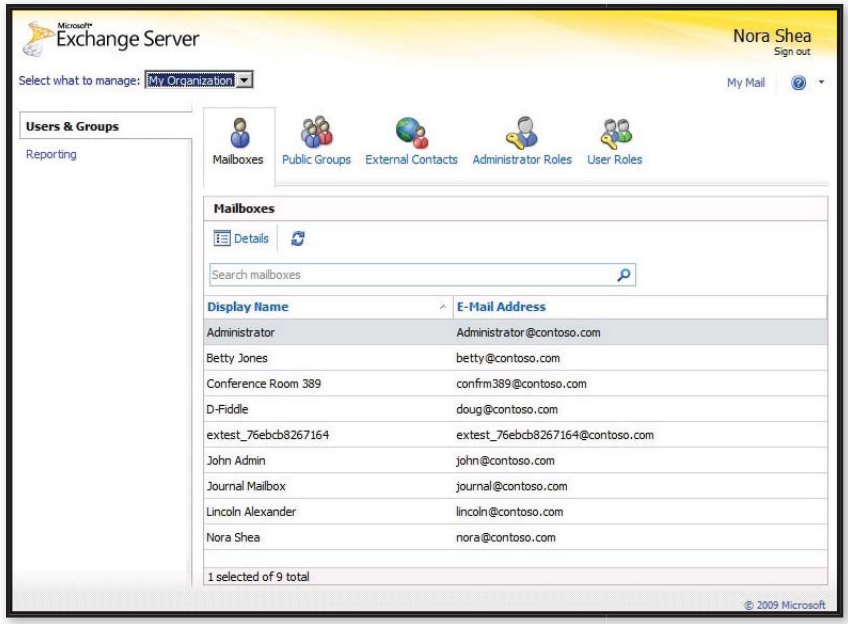

Figure 2: ECP UI with tabs for access to additional capabilities

you with only the EMS cmdlets and parameters that your roles allow you to have. For example, if you don't have the Mailbox Import Export role, you won't have access to the Export-Mailbox cmdlet.

In addition to the new virtual directories, Exchange 2010 adds a few new Windows services to the Client Access role that dramatically change how client connectivity works. The most significant of these new Windows services is the RPC Client Access (RPCCA) service, which moves the remote procedure call (RPC) endpoint for clients from the Mailbox server to the Client Access server. Clients talk to the Client Access server by using RPC, and the Client Access server makes the MAPI RPC connection on behalf of the user. The Client Access server uses the Active Manager client component to determine which Mailbox server hosts the user's database, then connects to that database, brokering the flow of MAPI traffic.

Outlook 2010 and Outlook 2007 clients can connect to RPCCA out of the box. However, if you're using Outlook 2003, you'll need to change the encryption setting. You can modify this setting on the Outlook client either by Group Policy or by making the change manually in the Outlook profile. Another option is to turn off the encryption requirement on your Client Access server, although this isn't what I would typically recommend.

By moving the MAPI endpoint for internal Outlook clients to the Client Access server, some new features are made possible. Perhaps the most anticipated features are online mailbox moves and an improved failover experience. Database failover is possible only when your Mailbox servers are in a database availability group (DAG). The Exchange Replication Service (MSExchangeRepl) monitors the databases and reports failures to the Active Manager component running on the Mailbox servers in the DAG. When the active copy of the database changes, MSExchangeRepl updates the Active Manager client component in the RPCCA service.

Because all Outlook clients now connect to the Client Access server directly, mailboxes can move between databases or even between Mailbox servers while the user is logged on and working in Outlook. The Client Access server now handles mailbox moves, and it does so through the Mailbox Replication Service (MRS). The process happens asynchronously, so you can create the move request, then continue working while the Client Access server does the heavy lifting. For more information about online mailbox moves, see "Moving Mailboxes the Exchange 2010 Way," on page 33.

If you move a mailbox to an Exchange 2010 database from an Exchange 2007 SP2 database or another Exchange 2010 database while users are online, there's no profile reconfiguration, no manipulation of DNS records, and no waiting for Active Directory (AD) replication. Because the user is connected to the Client Access server, only that server needs to be aware of the database location on the user's AD account. When the mailbox move request is initiated, the AD attributes are updated and the Client Access server in that site is made aware that a move request is in progress.

To fully enable the transition to RPCCA, another service also had to be moved out to the Client Access server. In Exchange, directory access is provided to clients through the Name Service Provider Interface (NSPI). In Exchange 2007 and earlier, referrals to an NSPI endpoint were provided by the DSProxy component. Exchange 2010 replaces DSProxy with the new Address Book service, which runs on the Client Access server. Moving the Address Book service and the RPCCA service gives you the freedom of no longer connecting to your Mailbox servers directly with Outlook.

Planning Your Client Access Infrastructure

Now let's take a look at some of the things you should do to adequately plan your Client Access infrastructure. Where to put the servers is likely a question at the top of your mind. The answer is the same in Exchange 2010 as it was in Exchange 2007: Everywhere! Client Access servers talk to Mailbox servers over RPC, so a Client Access server can establish a connection with a Mailbox server only in the same site. Because of this limitation, one of the requirements of Exchange 2010 is to have at least one server hosting the Client Access role in every AD site that includes Mailbox servers. This requirement isn't a problem if you're providing access to users only inside your network. But if you want users outside your network to access their email—through OWA, Outlook Anywhere, POP/IMAP, or ActiveSync—then you'll need to put more consideration into your Client Access server layout.

When deciding which Client Access servers to make Internet-facing, one of the important factors to consider is your external namespace. If you have more than one AD site with an Internet-facing Client Access server, it's difficult to pull off a single-namespace scenario. To understand why, let's look at two important parameters on your virtual directories, the InternalURL and the ExternalURL. The InternalURL contains the URL used by clients on the internal network. The ExternalURL contains the

URL that clients use over the Internet. For example, in OWA this might be https://mail.contoso.com/owa. If you have multiple Internet-facing Client Access servers, they'll each have their own ExternalURL defined for their virtual directories.

To illustrate the effect of these settings, let's see what happens when Lincoln, who has a mailbox in the Baltimore site on BAL-MB01, accesses a Client Access server in Seattle (SEA-CAS01). When Lincoln tries to use OWA on SEA-CAS01, the Client Access server realizes that Lincoln's mailbox is in Baltimore. If BAL-CAS01 has anything in the ExternalURL field for OWA, then SEA-CAS01 knows that BAL-CAS01 is Internet-facing and manually redirects Lincoln to use the ExternalURL for BAL-CAS01. However, if BAL-CAS01 doesn't have an ExternalURL, SEA-CAS01 uses the InternalURL defined on BAL-CAS01 to proxy the connection for the user. When proxied, if a user were to open a message in OWA, the request for the message would be passed from the user to SEA-CAS01 to BAL-CAS01 to the Mailbox server hosting the user's mailbox (BAL-MB01). Figure 3 illustrates this process.

You can see that when you have multiple Internet-facing Client Access servers, it makes sense for each site to have its own namespace. If they all have the same namespace, the ExternalURL might not resolve correctly to the site that hosts the user's mailbox. The other, more common, option is to have only one site with Internet-facing Client Access servers and have all users access it through a single namespace. If a user's mailbox isn't in that site, the connection is always proxied to the Client Access server that's in the right site.

Sizing the Servers

After you figure out where your Client Access servers should go, you'll need to decide how many to have and on what hardware they should run. I'll be the first to admit that sizing your Client Access servers is more of an art than a science. Unless you have a very simple environment, it's nearly impossible to tell ahead of time how many users will connect to a Client Access server, how frequently they'll connect, and what protocol they'll use. Yet, all of these factors need to go into the equation when deciding how many servers to have. Unfortunately, 10,000 users connecting through OWA doesn't generate the same load as 10,000 users using Outlook Anywhere or ActiveSync. To make it even more difficult, you need to consider many factors outside of the Client Access role that can lead to poor client connectivity experiences, such as TCP/IP connection limitations or a poorly sized AD implementation.

Microsoft has provided some general guidance for figuring out how many Client Access servers you need in "Understanding Exchange Performance" (technet.microsoft.com/library/dd351192.aspx). However, these numbers are generalizations and won't necessarily apply to specific implementations. The Microsoft guidelines for Exchange 2010 indicate that when a Client Access server is the only role on a server, you need three CPU cores in your Client Access server for every four CPU cores on your Mailbox servers in the same site. For example, if you have 2 Mailbox servers in a site with 8 cores each, for a total of 16 cores, you should ensure that the total number of Client Access cores in that site adds up to at least 12, which could mean 3 servers with 4 cores each, or even 6 servers with 2 cores each.

For memory in Client Access servers, you need a minimum of 4GB plus 2GB for each CPU core, up to a maximum of 16GB. For example, if you have a Client Access server with 4 CPU cores, you need 8GB of memory in the server. Start with these guidelines, but also do the research and testing needed to ensure that your Client Access infrastructure is properly scaled.

What do I mean by doing the research and testing? To start, you need to understand what protocols your clients are using to access the Client Access server. Your client traffic is going to be composed of MAPI, Outlook Anywhere, ActiveSync, EWS, and OWA. If you're currently using Exchange 2007, you can start gathering this data with some simple performance monitoring on your existing Mailbox and Client Access servers. There are performance counters available for each of these protocols, and capturing this data over a period of days, weeks, or months can give you a good idea of which protocols clients use to access mail. The longer you can collect data, the more accurate your estimates will be. In addition to Performance Monitor, you can use tools such as the Exchange Server User Monitor (www.microsoft.com/downloads/details.aspx?FamilyID=9a49c22e-e0c7-4b7c-acef-729d48af7bc9), which helps you analyze MAPI traffic.

After you have some internal statistics, you can use the data to determine how to size your servers. You should deploy your server's hardware configuration in a lab environment, then test it by generating the load and protocol mix that you estimated in your research. The 2010 version of Exchange Load Generator is currently in beta, but you can use this tool to simulate the client protocol traffic (www.microsoft.com/downloads/details.aspx?familyid=CF464BE7-7E52-48CD-B852-CCFC915B29EF). With simulation testing, you can determine what your hardware bottlenecks are and decide to scale up the hardware configuration or throw more Client Access servers in your array if necessary.

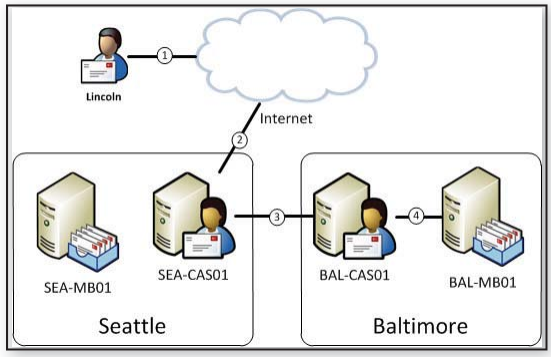

Figure 3: A message request proxied through the InternalURL

Ready to Roll

As you can see, the Client Access server role is integral to any Exchange 2010 or Exchange 2007 deployment, and getting the details right will ensure efficient mail transfer. Now that you have a better understanding of the Client Access server and its increased importance in Exchange 2010, you're ready to look at some more specific guidance for deploying the role and getting the most out of it—which I'll address in my next article in this series.

InstantDoc ID 125061

Ken St. Cyr
(ken.stcyr@microsoft.com) is a solution architect at Microsoft with more than 10 years of industry experience. He's a Microsoft Certified Master in Directory Services and the author of *Exchange Server 2010 Administration Instant Reference* (Sybex).

AppLocker in Windows Server 2008 R2 and Windows 7

Whitelisting technology provides application access control

by Jan De Clercq

Windows Server 2008 and Windows Vista include the BitLocker feature, which provides volume-level encryption for bits stored on Server 2008 and Vista computers. Windows Server 2008 R2 and Windows 7 include AppLocker, which is an application locker that lets Windows administrators provide application access control to restrict which applications can run on their domain's workstations and servers.

AppLocker is an enhanced version of Windows Server 2003's Software Restriction Policies. SRPs and AppLocker tackle the problem of execution of malicious code on Windows platforms. You can also use SRPs and AppLocker to block user access to games such as Minesweeper, or to prohibit the startup of a browser that isn't standard in your organization.

In this article I compare AppLocker with SRPs, concentrating on the caveats and pitfalls you must be aware of in evaluating AppLocker for use in your environment. Table 1 lists the main differences between SRPs and AppLocker.

Blacklisting and Whitelisting

Malware-protection programs such as antivirus and antispyware software often use a technique referred to as *blacklisting* to protect computers. Programs that employ blacklisting allow everything to be stored on a computer other than files that are infected with threats listed on the blacklist. If a file is infected, these programs will either delete or quarantine it.

Table 1: SRPs vs. AppLocker

Feature	Software Restriction Policies (SRPs)	AppLocker
Default policy	Blacklisting (unrestricted security level)	Whitelisting
Rule conditions provided	File hash Path Certificate Network zone (Internet zone)	File hash Path Publisher
Rule types provided	Defined by security level: Disallowed Basic user Unrestricted	Allow Deny
Rule scope	Applies to all users	Can be set for specific user or group account
Supports audit-only mode	No	Yes
Includes wizard to automatically create a whitelist	No	Yes
Supports rule import and export	No	Yes
Supports PowerShell	No	Yes

An emerging approach to combating malware is *whitelisting*. Whitelisting takes an opposite approach to blacklisting—that is, the protection program blocks everything except the files that are on its whitelist.

When it comes to protecting computers against the execution of unwanted malware, whitelisting is preferable to blacklisting. Whitelisting eases the life of the administrator, because in today's interconnected world, users are typically *allowed* to run fewer applications than they should be *blocked* from running—including an almost unlimited number of unknown malicious executables that users might download from the Internet. However, whitelisting creates the risk of locking yourself out if you don't use it properly. For example, you might neglect to add your management applications to the whitelist. In addition, you can inadvertently prevent your users from working if you forget to add one of their applications to the whitelist.

SRPs and AppLocker both support whitelisting and blacklisting, although they have different default policies. AppLocker uses whitelisting by default, thereby blocking everything; the administrator must explicitly define the applications that can run. The default SRP configuration uses blacklisting, which allows all applications to run; the administrator must define exceptions for any applications to be blocked. Setting up whitelisting with SRPs is difficult, which is why most admins use it only for blacklisting applications. AppLocker is much better suited to provide whitelisting-based protection for controlling applications.

Setting Up AppLocker Rules

To begin, you must know how to configure application restriction rules in AppLocker. As with SRPs, you can use Group Policy Object (GPO) settings to configure and enforce AppLocker rules. You can also use PowerShell cmdlets to configure AppLocker rules (this option isn't available for SRPs). (For information about using PowerShell cmdlets with AppLocker, see the MSDN Windows PowerShell Blog entry "Getting Started with AppLocker management using Powershell" at blogs.msdn.com/powershell/archive/2009/06/02/getting-started-with-applocker-management-using-powershell.aspx.)

As Figure 1 shows, the AppLocker GPO is in the \Computer Configuration\Windows Settings\Security Settings\Application Control Policies container. Notice that you can also configure rules in the Software Restriction Policies container. The two technologies can coexist—you can leverage SRPs on all Windows platforms after Windows XP and Windows Server 2003, but AppLocker is available only on Server 2008 R2 and Windows 7 Ultimate and Enterprise. Because the policy rules that the two technologies use are so different, Microsoft doesn't provide an automatic conversion from SRP to AppLocker policies (e.g., if you upgrade a Server 2008 machine to Server 2008 R2).

AppLocker supports three rule types: Executable Rules, Windows Installer Rules, and Script Rules. These rule types are grouped in rule collections and appear as subcontainers of the AppLocker container in the GPO settings, as Figure 1 shows.

- Executable Rules can allow or prevent *.exe and *.com files from running.
- Windows Installer Rules can allow or prevent the execution of *.msi (Windows Installer) and *.msp (Windows Installer patching) files.
- Script Rules can allow or prevent the execution of different script file types (*.ps1, *.bat, *.cmd, *.vbs, *.js).

When you right-click one of the three rule collection containers, AppLocker gives you three options for creating rules: Create New Rule, Automatically Generate Rules, and Create Default Rules.

Create Default Rules. The preferred option for getting started with AppLocker rule definitions is Create Default Rules. Default rules are generated automatically; these rules are tailored to let Windows run and to let you do your administrative work—both of which are important, considering AppLocker's default whitelisting approach and the risk of locking yourself out. As a safety net, AppLocker prompts you to automatically create the default rules if you try to create a new rule and haven't yet created the default rules.

AppLocker's default rules are relatively open. For example, they include a rule that gives members of the local administrators group access

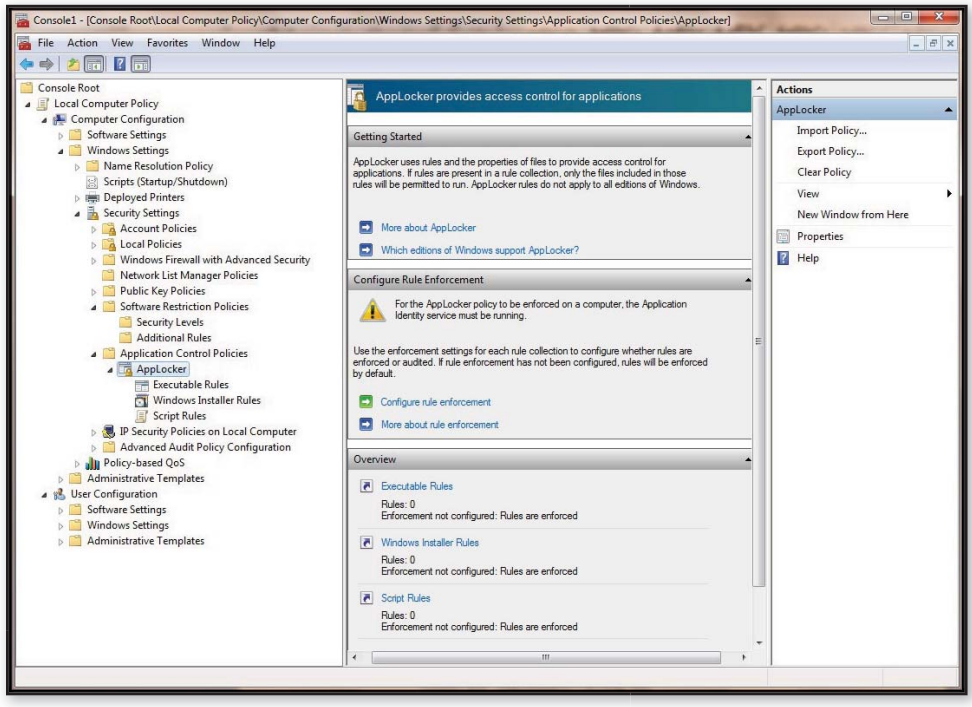

Figure 1: Accessing the AppLocker GPO

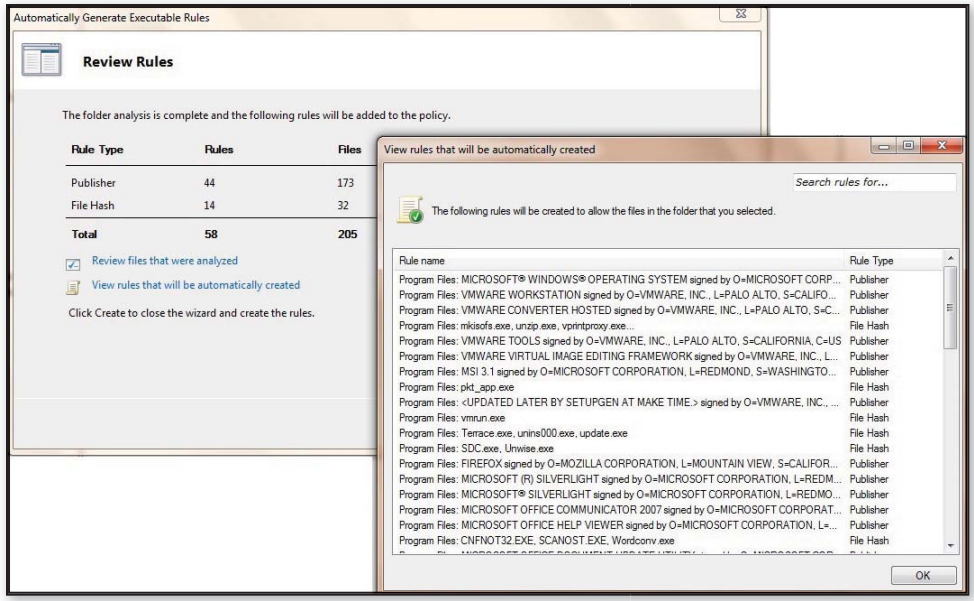

Figure 2: Automatically generating an AppLocker rules list

to all local files. An AppLocker best practice is to first create default rules, then refine them using more restrictive rules that you create manually through the Create New Rule option (which I explain later). Default rules can be created separately for each of the three rule types: Executable Rules, Windows Installer Rules, and Script Rules.

Automatically Generate Rules. With the Automatically Generate Rules option, AppLocker basically generates a whitelist for you. Based on the file folder you provide in the automatic rule generation wizard, AppLocker will propose a set of rules for the files in that particular folder. This important new AppLocker feature isn't included in SRPs. With SRPs you must define the whitelist yourself. To create an AppLocker whitelist for a certain category of machines, I recommend that you use a reference computer. Sharing such a whitelist with other computers and importing the whitelist into a GPO is relatively simple thanks to AppLocker's easy-to-use export/import mechanism.

To automatically generate a rule set, select Automatically Generate Rules from the Executable Rules, Windows Installer Rules, or Script Rules context menu. On the wizard's first screen, select a file system location on the reference machine, indicate the users or groups you want the whitelist to apply to (an important option that isn't available in SRPs), and provide a name for the resulting rule set.

Next, select whether you want to create rules based on the file's hash or path. Using the hash thumbprint is the most unique way to identify a file, but it has the disadvantage of requiring revisions of the rule when the file is updated (e.g., after a patch cycle). In addition, using hashes can negatively affect system performance.

The wizard also gives you the option of reviewing the list of files that it analyzed (including cancelling creation of a rule for a given file), previewing the rules list, and searching the list. Figure 2 shows the final step of actually creating the rules list.

Create New Rule. Another approach to creating AppLocker rules is to define them manually, through the Create New Rule option. You'll typically use this option as a last step to refine the rules AppLocker created for you using one of the options that I discussed previously.

From the rule creation wizard, you can create new allow or deny rules for files, select the group or user you want the rule to apply to, and choose whether you want AppLocker to identify the files using the file's hash, the file path (file

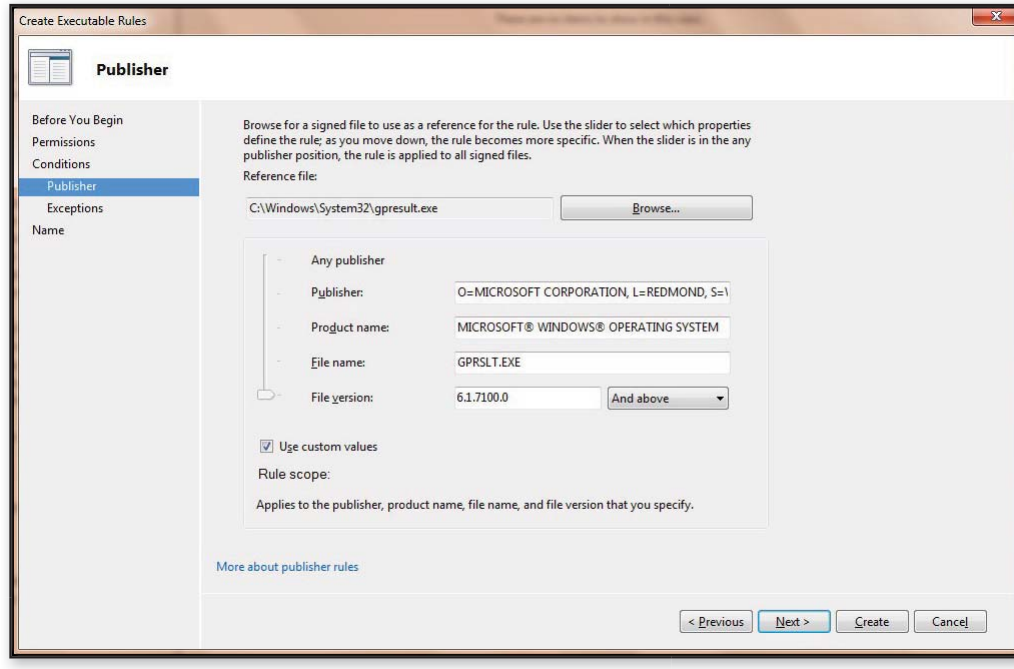

Figure 3: Identifying files by publisher

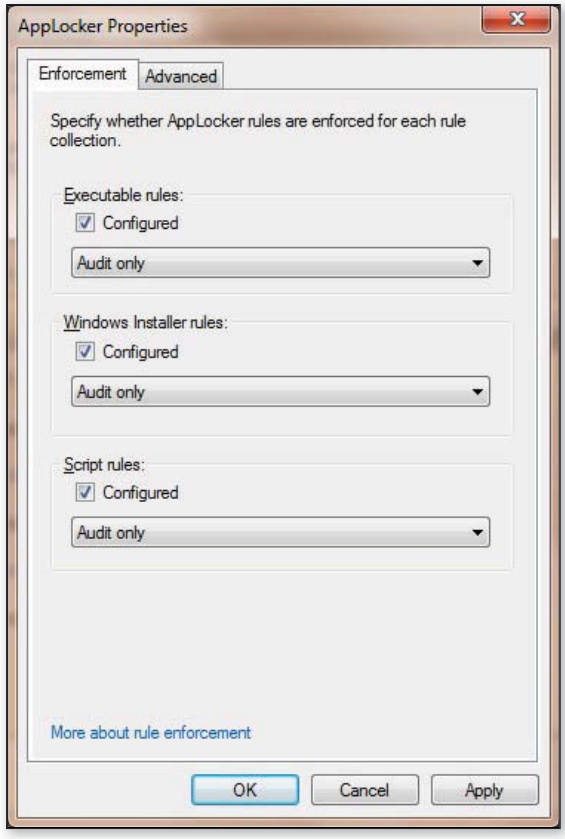

Figure 4: Setting AppLocker enforcement options

system location), or the file's publisher. AppLocker refers to these identification options as *primary conditions.* You can then further narrow down the identification of files by specifying exceptions to the primary condition in the wizard (hash-based, path-based, or publisher-based).

If you choose publisher-based identification, AppLocker will identify the file through the file's digital signature that was applied by the publisher as part of its software signing process. Figure 3 shows the options for publisher-based identification. You can use the slider bar to limit the identification of files to only the publisher name (most general) or expand identification as far as the exact file version (most specific). The wizard also lets you specify custom values for the publisher, product name, file name, and file version fields.

Note that AppLocker doesn't offer the network zone (aka Internet zone) file identification option that SRPs provide, which lets you use the Internet zone of the website from which code was downloaded to identify the code.

Enforcing AppLocker Rules

Like SRPs, AppLocker isn't enabled by default. Even when you're done creating rules, AppLocker won't immediately enforce the rules on your clients. Rule enforcement requires two additional steps. First, you must specify whether you want to enforce your rules or run them only for auditing purposes. Second, you must ensure that the Application Identity Service is running on the targeted machines.

You can set AppLocker's enforcement options in the AppLocker GPO's properties. As Figure 4 shows, you can specify whether a rule is configured (the default is *not* configured) and indicate whether the rule should be enforced or run only in audit mode, for each of the three main rule collections (i.e., Executable rules, Windows Installer rules, and Script rules).

The *Audit only* option is a useful new feature that isn't available with SRPs. When a rule collection is set to *Audit only* mode, the rules within that rule collection aren't enforced, but any time a user runs an application that's affected by a rule, information about the rule and the application write to the local machine's AppLocker event log container.

I recommend that you select the Configured check box and choose *Audit only* in the drop-down menu for each of the three rule collections. Not only does this protect against locking yourself out, but it also lets you see whether your rules are catching the correct applications, as well as whether they're too permissive or too restrictive. I also recommend that you use the *Audit only* mode until you've recorded and evaluated all the rules' effects and side effects.

Note that the Advanced tab of the AppLocker container's properties refers to a fourth AppLocker rule collection: DLLs, to cover the *.dll and *.ocx file formats. Microsoft set this rule collection apart in the Advanced tab because of the performance impact DLL checking has when it's enabled. In addition, the process of whitelisting all the allowed DLLs creates a significant amount of administrative overhead. You should enable AppLocker DLL protection only in organizations with extremely critical IT security (e.g., government or defense organizations).

The last step in guaranteeing AppLocker enforcement is to make sure the Application Identity Service is enabled on your Server 2008 R2 and Windows 7 machines. This service is set to manual startup by default. To properly use AppLocker, you must set the service to start up automatically. You can use GPO settings to configure all your machines at once. Because anyone with local administrator rights can stop the service and therefore bypass AppLocker policy enforcement, you need to keep tight control over your administrator accounts.

A Major Step Forward

Like SRPs, AppLocker requires regular rule updates to properly deal with patches and new versions of protected applications. AppLocker can't yet deal with software updates in a dynamic and silent fashion. For this purpose, certain third-party whitelisting applications (e.g., Coretrace's Bouncer, Bit9's Parity) will perform better. In addition, these applications provide broader platform and file-type support. However, AppLocker is a major step forward for application whitelisting in Server 2008 R2 and Windows 7, compared with SRP blacklisting. Windows administrators will appreciate AppLocker's ability to automatically create whitelists, to run in audit-only mode, and to limit rule application based on user and group accounts.

InstantDoc ID 104625

Jan De Clercq
(jan.declercq@hp.com) is a member of HP's International Expertise Team and focuses on architecture for Microsoft-based IT infrastructures, identity management, and security. He's co-author of *Microsoft Windows Security Fundamentals* (Digital Press).

The SCOM Service Level Dashboard

Follow these steps and know your uptime at a glance

by Richard Raseley

The Service Level Dashboard (SLD) Solution Accelerator from Microsoft works with System Center Operations Manager (SCOM) 2007 R2 to assist both technical and non-technical members of your organization in monitoring application and service availability and network and system performance. The primary benefit of the SLD is the ability to answer at a glance the frequent question, "How are we doing?"

The dashboard, which Figure 1 shows, is presented via a Microsoft Office SharePoint site that can be exhibited on a large display unit in your environment or sent as a link to stakeholders in your organization. It also provides a rich set of data that allows for advanced reporting of issues and trends, helping implementers and architects to prove the effectiveness of initiatives or acting as a source of data to prove the need for new initiatives.

Design Concepts and Prerequisites

The design for my scenario (as Figure 2 shows) includes three servers and any number of monitored applications, clients, or servers. All monitored clients report directly to the SCOM server, which in turn works with Microsoft SQL Server to manage and store the data. Once the SharePoint Server instance is configured with the SLD components, it pulls any relevant data (depending on the metrics you've selected to monitor) and displays it in the dashboard site.

The SLD accelerator requires:

1. SCOM 2007 R2 with Data Warehouse and Reporting components installed
2. SharePoint Server 2007 SP1 or Windows SharePoint Services (WSS) 3.0 SP1
3. SQL Server 2005 SP2 or SQL Server 2008 (SCOM and WSS can be installed on the same server, but it's recommended that you at least have a separate SQL Server instance—if not an entire server, depending on the network size—for the SCOM databases)
4. Microsoft .NET Framework 3.5 (all servers)

SLD Installation

Once you've verified you meet all the prerequisites, you can begin the SLD installation process by installing the SLD management pack. SCOM uses management packs to extend its ability to monitor applications, products, and services that weren't initially built into the application. This approach allows for a high level of flexibility and customization of your environment. The SLD package contains one of these management packs, which contains the appropriate data to allow the SLD site to pull the necessary data from the SCOM server.

You need to download the file *Service Level Dashboard 2.0.zip* from bit.ly/bxrKMD and extract it to a location that the SCOM server can access. Once you've extracted the files from the zip, log on to the

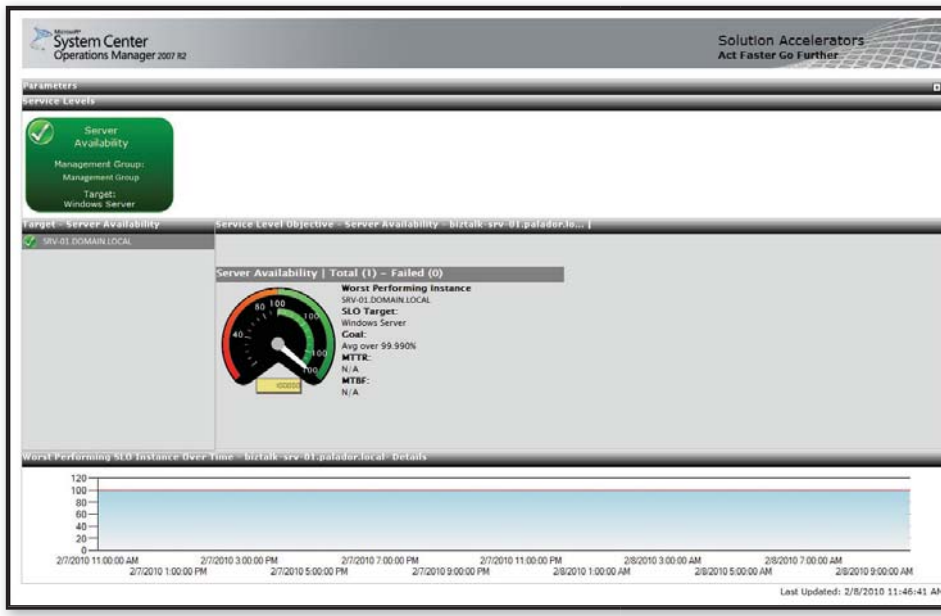

Figure 1: The Service Level Dashboard Solution Accelerator

server running SCOM with an account that has SCOM administrator privileges and launch the SCOM Console. Click Administration in the left pane of the console, then Import Management Packs in the right pane of the Administration area. In the Import Management Packs window in the Select Management Packs screen, click Add, then Add From Disk in the drop-down menu. From the *Select Management Packs to Import* window, navigate to the location where you extracted the zip file, select Microsoft.EnterpriseServiceMonitoring.ServiceLevelDashboard.R2.mp, and click Import. This will take you to the Import Management Packs window. Once you're there, click Install.

After you've imported the SLD management pack into your SCOM installation, you can install the SharePoint components. Again, log on to the server running SCOM with an account that has SCOM administrator privileges and navigate to the folder where you extracted the zip file. Execute the SLD installation file that matches your processor architecture—ServiceLevelDashboard_x64.msi if your system has a 64-bit OS or ServiceLevelDashboard_x86.msi if your system has a 32-bit OS.

On the *Welcome to the Service Level Dashboard 2.0 Setup Wizard* and *End-User License Agreement* screens, click Next. In the Operations Manager 2007 R2 Information window, which Figure 3 shows, enter the Application Pool Identity (the user under whose authority you want the SLD to run), the Operations Manager Data Warehouse Server Name (the name of the SQL Server where you installed the Operations Manager Data Warehouse components), and the name of the Operations Manager Data Warehouse Database (OperationsManagerDW by default).

In the Windows SharePoint Services 3.0 Information window, enter the Site Owner Login (the account of the individual that you want to assign the site owner role to), the Site Owner Email Address, the SharePoint Database Server Name (the fully qualified domain name of the server that hosts the databases for the SharePoint instance), and the Service Level Dashboard URL (the URL and port that you want to be associated with the initial SLD site—this can be modified later). Note that if you specify a URL other than the name of the server that's hosting the site, you have to create the appropriate records in your DNS infrastructure. When you're done with that, just click Next on the next three screens of the wizard.

SLD Configuration

The Service Level Tracking feature allows organizations to define Service Level Objectives that are used to monitor the availability and health of applications, services, and systems. For this example, I'll create a Service Level Objective that will monitor the availability of a server, with a target of 99.99 percent uptime.

First, log on to the server running SCOM with an account that has local administrator privileges and launch the SCOM Console. Click Authoring, then Management Pack Objects, Service Level Tracking, then

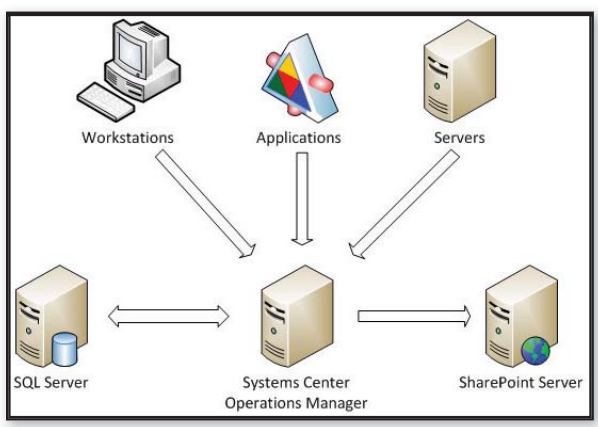

Figure 2: Example scenario layout

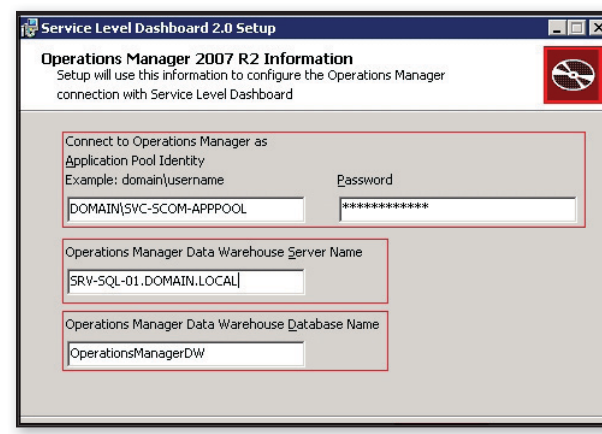

Figure 3: The setup screen for Operations Manager

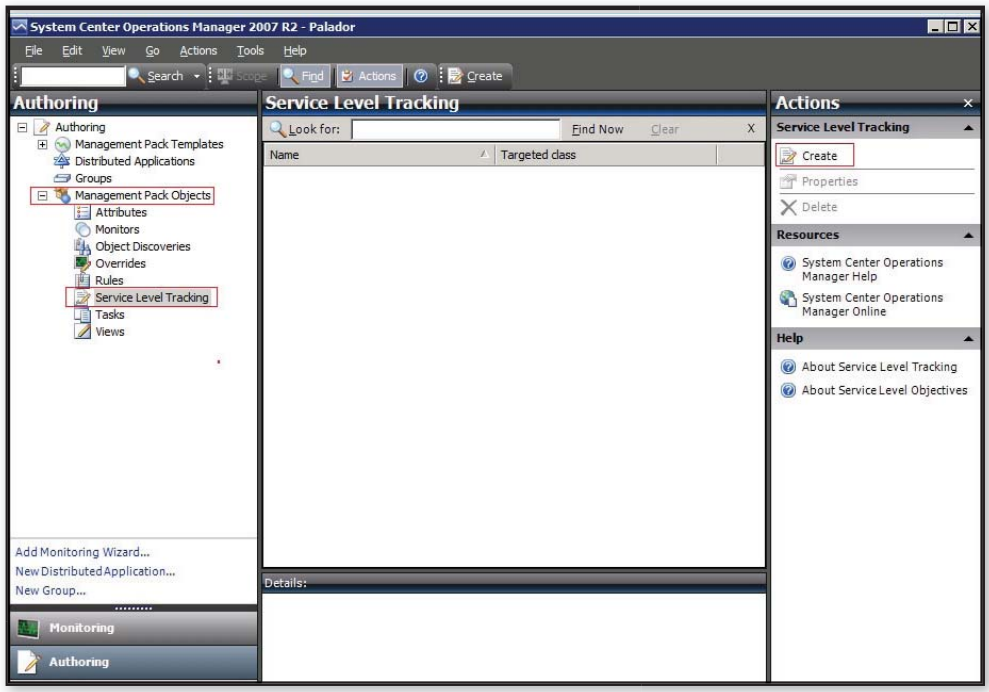

Figure 4: Finding Service Level Tracking

Create in the right pane, as Figure 4 shows. In the left pane of the Authoring area, expand Management Pack Objects, click Service Level Tracking, then click Create in the right pane. Input a name for the new Service Level Objective (e.g., Server Availability) and a description in the Service Level Tracking window and click Next.

Click the Select button in the Targeted Class section of the Objects to Track window. In the Select a Target Class window, which Figure 5 shows, select the application, class, or group of objects you want to monitor. I'll select Windows Server. In the Service Level Objectives window, click Add, then Monitor State SLO. In the Service Level Objective Monitor State window, enter a name for the Service Level Objective (e.g., Server Availability), enter a Service Level Objective Goal (%) of 99.99, and click OK, then Next. Review your selections for accuracy In the Summary window and click Finish, then Close.

The primary benefit of the SLD is the ability to answer at a glance the frequent question, "How are we doing?"

Your Service Level Objective has been defined, so you can configure your SLD site.

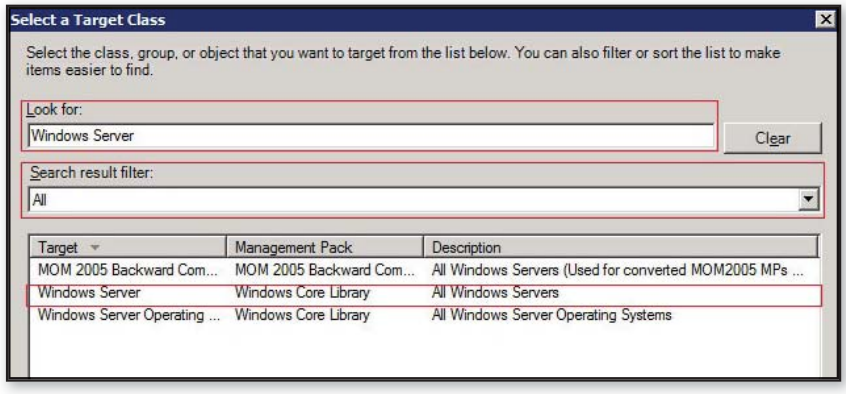

Figure 5: The Select a Target Class window

Log on to any domain member with the user account that was specified in the SLD Installation section (e.g., DOMAIN\Administrator). Launch Internet Explorer and navigate to the URL that was specified in the SLD Installation section (e.g., http://ServiceLevel Dashboard:80). Then, from the main page, click on the Site Actions drop-down menu and choose Edit Page.

In edit mode, you can see the hidden Dashboard Configuration web part. This web part lets you choose which service level objectives are displayed on your SLD site. In the web part, check Server Availability and click Apply Filter. Also on this page, note the Dashboard Refresh, Dashboard Default View, and Aggregation Type options. These control how the service level objective data is displayed and updated on the site. Click the Exit Edit Mode link to confirm that the service level objective has been added to the dashboard site.

You now have a complete dashboard site with a monitor that will tell you if you're meeting your uptime target of 99.99 percent. Using this at-a-glance view of the health of your infrastructure can let you instantly address your stakeholders' fears and worries. It'll free up your IT department's time to ensure that service levels remain high, reducing the need for firefighting. The SLD Solution Accelerator provides this information by seamlessly integrating into your existing infrastructure, taking advantage of SCOM, SharePoint, and SQL Server. Deployment is easy, and it's a solid investment in your monitoring infrastructure.

InstantDoc ID 125078

Richard Raseley
(Richard@Palador.com) is a systems engineer with Palador in Seattle. He focuses primarily on design and deployment of Microsoft technologies, including Active Directory, Exchange Server, Hyper-V, SharePoint, and Microsoft's server line of OSs.

WinConnections ...

Providing the vision

THE CONVERSATI

WINDOWS CONNECTIONS

SharePoint CONNECTIONS

MICROSOFT EXCHANGE CONNECTIONS

UNIFIED COMMUNICATIONS CONNECTIONS

NOVEMBER 1-4, 2010

MANDALAY BAY
RESORT & CASINO

LAS VEGAS, NEVADA

EARLY BIRD DISCOUNT! *Register by July 29 and book a minimum of 3 nights at Mandalay Bay and you'll receive a $100 Mandalay Bay Gift Certificate.*

WINDOWS CONNECTIONS //

- Featuring the industry's most well-known and respected technology experts
- Learn to "do more with less" while increasing your IT skills
- Original content you won't find anywhere else
- Opportunities to make new connections amongst your peers
- Connect with key technology solution vendors

FIVE KEY FOCUS AREAS:

- Business Technology
- Building Your Skill Set
- Windows Server 2008 R2
- Windows 7
- Virtualization

Exchange Connections & Unified Communications //

Exchange and OCS Solutions for the Real World:

- Deployment
- Management
- Maintenence
- Microsoft Business Productivity Online Services (BPOS)
- Information Protection
- Useful Features in Service Pack 1
- Integration of Exchange with SharePoint (and other collaboration solutions)
- Best ways to use Unified Communications in your organization
- Sessions that Cover Exchange 2003 and Exchange 2007 and preparing for Exchange 2010

SharePoint Connections //

- Upgrade and Deployment to SharePoint 2010
- Enterprise & Web Content Management
- LINQ
- Business Connectivity Services
- Silverlight
- Workflow
- Virtualization
- Claims-based Authentication
- Enterprise Search
- Business Intelligence
- Security
- SharePoint Connections Bonus! No Code Solutions Track

Microsoft SharePointPro CONNECTIONS SQL

intelligence
to keep you and your company
competitive in today's market!

ON BEGINS HERE

ENJOY A PREMIERE LAS VEGAS HOTEL!

Mandalay Bay Resort And Casino

The Mandalay Bay Resort and Casino offers elegance, excitement, and escape. Indulge in award-winning restaurants helmed by celebrity chefs, an enormous beach-front pool, the luxurious spa, and top-notch entertainment. You'll find plenty of ways to relax with your colleagues from around the world at the end of each of our knowledge-packed days.

Only Microsoft and Industry Experts speak at WinConnections!

A sampling of our speakers ...

TED PATTISON
CRITICAL PATH TRAINING, LLC

STEVE FOX
MICROSOFT

DAN HOLME
INTELLIEM, INC.

MICHAEL NOEL
CONVERGENT COMPUTING

ANDREW CONNELL
CRITICAL PATH TRAINING, LLC

TODD BAGINSKI
INTELLIGENT EFFECTS

DON JONES
CONCENTRATED TECHNOLOGY

KEVIN LAAHS
HP

RHONDA LAYFIELD
CONSULTANT/TRAINER

GREG SHIELDS
CONCENTRATED TECHNOLOGY

MARK MINASI
MR&D

TONY REDMOND
TONY REDMOND AND ASSOCIATES

KIERAN MCCORRY
HP

CHRIS AVIS
MICROSOFT

CHECK WEB SITE FOR DESCRIPTIONS OF SESSIONS AND WORKSHOPS

www.WinConnections.com • 800.505.1201 • 203.400.6121 • Register Today!

SERVER

Windows IT Pro

TECH Conferences
PENTON MEDIA

Hyper-V Live Migration FAQs

Questions answered and misconceptions set straight

by Michael Otey

Without a doubt, the introduction of Live Migration with Windows Server 2008 R2 was that release's most important virtualization feature. Live Migration lets you move a running virtual machine (VM) between Hyper-V hosts with no downtime. However, as with all new technologies, there are some questions and misconceptions about this new feature. Let's tackle some of the frequently asked questions about Live Migration.

Q: Is the new Cluster Shared Volumes (CSV) feature required for Live Migration?

A: No. Contrary to a popular misconception, Live Migration doesn't require the use of CSV. As Microsoft likes to phrase it, CSV facilitates Live Migration. It essentially allows multiple Hyper-V VMs to access the same set of VM files located on shared storage—typically a LUN on a local SAN. That makes it easy to set up Live Migration, because you can simply give each VM access to the shared storage. Performing live migrations with CSV is fast because both VMs can access the same stored files and only the memory needs to be moved between the Hyper-V hosts. Performing live migrations without CSV takes a bit longer because the storage must be moved between the nodes.

Q: Will Live Migration work with Linux VMs?

A: Yes. Live Migration works at the Hyper-V host level, so it's independent of the guest OS running in the VM. Live Migration works just as well with Linux-based Hyper-V guests as it does with Windows-based guests. In the lab, I've successfully used Live Migration with both the Ubuntu and openSUSE Linux distributions.

Q: Do you need to use Microsoft System Center Virtual Machine Manager (VMM) with Live Migration?

A: No. You don't need to use VMM, which is a great tool for managing multiple Hyper-V (and even VMware ESX Server) hosts and their respective VMs. You can use Server 2008 R2's Failover Cluster Manager with Live Migration. To do so, expand the cluster containing the Hyper-V nodes. Then expand Nodes and right-click the VM node you want to work with. Select the *Live migrate virtual machine to another node* option from the pop-up menu.

Q: How does Live Migration handle different processor architectures?

A: It's important to realize that for Live Migration to work, the physical processors in the Hyper-V host servers need to be compatible. You can't perform a live migration between Hyper-V hosts that utilize different processor manufacturers. In other words, to perform a live migration, both Hyper-V hosts must have Intel processors or both must have AMD processors—you can't perform a live migration if one Hyper-V host has an Intel processor and the other Hyper-V host has an AMD processor.

The processors in the different Hyper-V hosts don't have to be identical, however. For example, one Hyper-V host can have an Intel Core 2 E6400 running at 2.13GHz and the other Hyper-V host might have an Intel Core 2 E8500 running at 3.16GHz. A Hyper-V feature called *processor compatibility mode* lets you perform live migrations between Hyper-V hosts with different processor families. The processor compatibility mode is turned off by default, but you can turn it on by opening the Hyper-V Manager and accessing the VM's CPU properties. Select the *Migrate to a physical computer with a different processor* check box. You can learn more about Hyper-V's processor compatibility in the Microsoft white paper "Windows Server 2008 R2 Virtual Machine Processor Compatibility Mode," which you can download by going to download.microsoft.com/download/F/2/1/F2146213-4AC0-4C50-B69A-12428FF0B077/VM%20processor%20compatibility%20mode.doc.

InstantDoc ID 103640

Michael Otey
(motey@windowsitpro.com) is the technical director for *Windows IT Pro* and *SQL Server Magazine* and coauthor of *Microsoft SQL Server 2008 New Features* (Osborne/McGraw-Hill).

Going Virtual with SharePoint 2010

What you need to know to get it right

Server virtualization technologies have become so commonplace that they're the de facto standard for server deployment in many organizations. It's becoming more and more common to run into data center environments that operate with the assumption that all new servers will be deployed as virtual machines (VMs), unless there's some specific reason not to virtualize. This is a significant change from even just a few years ago when the situation was reversed and servers were deployed on physical equipment unless there was a specific reason to virtualize.

So, what about Microsoft SharePoint? Should you virtualize some or all of a SharePoint environment and take advantage of the consolidation, optimization, and flexibility options that a virtualization infrastructure provides? The reality is that SharePoint environments, particularly those running SharePoint Server 2010, can be robustly deployed on virtual servers as long as sufficient resources are allocated to virtual guests and the virtual hosts are scaled correctly. Deploying SharePoint improperly in a virtual environment can lead to slowness and other performance problems, and can decrease management's confidence in virtualization as a whole. So, before you virtualize SharePoint, it's vital that you know the requirements and design criteria for both the virtualization technology and SharePoint.

Virtualization Infrastructure Requirements and Recommendations

The key to a stable and high-performance virtualized SharePoint environment is using the proper architecture in the virtualization hosts. Out-of-the-box settings and slow disks might work for a test environment, but specific requirements need to be met when building the host system for proper performance to be achieved. Therefore, be sure to follow these minimum requirements when you design the virtualization host infrastructure:

- The processors must support hardware-assisted virtualization, which is available in processors that include a virtualization option. Specifically, this means processors with Intel Virtualization Technology (Intel VT) or AMD Virtualization (AMD-V) technology.
- Hardware-enforced data execution prevention (DEP) must be available and enabled.
- SharePoint guests must be deployed on a Hyper-V hypervisor or a third-party hypervisor that's part of the Server Virtualization Validation Program. (For information about this program, see the sidebar "Microsoft's Virtualization Support Story.")
- Sufficient memory must be allocated for the host OS. If you're using Hyper-V, you need to reserve at least 1GB of RAM for use by the Hyper-V host. If you're using a third-party hypervisor, check with the individual provider to determine the minimum amount of memory required.
- A dedicated NIC must be allocated for host management. This NIC must be separate from the NICs used by the VMs.
- Use multiple storage arrays or sets of disk spindles. Best practice is to allocate a dedicated storage array or set of disk spindles for the host OS, another for the guest OS, and at least two more for logs and database volumes in virtualized SQL Server sessions.
- Fixed-size or pass-through Virtual Hard Disks (VHDs) must be used. All VHDs used by SharePoint servers need to be either fixed-size or pass-through (raw) disks that are directly connected to a volume on the host storage. Pass-through disks give you the fastest performance, which is highly recommended for SharePoint servers. Fixed-size disks are faster than dynamically expanding disks, which can suffer performance hits when they're resizing.
- A 2:1 ratio for the number of virtual processors to physical cores must be used. A virtual host that has too many allocated virtual CPUs can be overloaded and perform poorly. Therefore, you need to have a 2:1 ratio (or less) for the virtual processor to physical core ratio. For example, if your host is a 2-processor quad-core system (8 cores total), the maximum number of virtual processors that can be allocated and running at any one time is 16. If each VM is allocated 4 virtual processors, the number of running VMs is capped at 4 on that host.

In addition to these technical requirements for the virtualization host, you need to keep in mind these recommendations when you design your virtual environment:

- You should allocate a dedicated NIC for failover. If you're using virtual host failover software such as Hyper-V Live Migration, you should use a dedicated NIC for the failover.
- You should give as much memory and as many processor cores to your virtual hosts as your

Michael Noel
is a partner at Convergent Computing, a Microsoft MVP, and the author of books on SharePoint, ISA Server, and Exchange Server. His latest book is *Windows Server 2008 Unleashed* (Sams).

budget allows. Virtual hosts with multiple multicore processors and large amounts of RAM (64GB or more) are becoming commonplace because of the virtual host software's ability to take advantage of the additional resources and because host failover solutions require additional resources. When it comes to sizing virtualization hosts, there's a sweet spot that balances the cost of the additional components against the need to have fewer hosts. Generally, the virtualization overhead required to run virtual servers is only 5 percent, so the cost of adding memory and processor cores is more than made up by the advantages of having those additional resources.

- You should run only the virtualization software and the virtualization role on the virtual hosts. (The two exceptions are antivirus and backup software.) Overloading a virtual host with other software or other server roles can significantly degrade guest performance. In addition, from a Windows Server licensing perspective, running any roles other than the virtualization role on a Windows server requires one additional license. However, if the host runs only virtualization host software, the host OS isn't counted when determining the number of Windows licenses that are used as part of Microsoft's virtualization licensing program.
- You shouldn't install all the SharePoint roles and the SQL Server role on the same VM for performance reasons. Even small environments should use at least two VMs—one for the SQL Server database role and one for the SharePoint front-end and application roles.

Software Recommendations and Licensing Notes

It's highly recommended that you use the latest virtualization host software from your particular vendor. For example, the latest version of Hyper-V is included with Windows Server 2008 R2. Hyper-V 2.0 has significant performance improvements over Hyper-V 1.0, such as I/O improvements for fixed-size VHDs. Hyper-V 2.0 also has new features such as Core Parking, Live Migration, TCP Offload, Jumbo Frames, and support for Second-Level Address Translation (SLAT)–enabled processors. If you're virtualizing SharePoint on Hyper-V, you should also consider deploying the virtual host on Server Core to minimize its security footprint, OS disk overhead (2GB versus 10GB), and memory use.

If you're managing multiple virtual host machines, centralized management software is also recommended. For example, Microsoft offers System Center Virtual Machine Manager (VMM) 2008 R2 for virtualization management. It allows for physical to virtual (P2V) server migration, server template libraries, and management

Microsoft's Virtualization Support Story

The support story for Microsoft products running on virtualization hardware is long and complicated. Until several years ago, Microsoft offered limited support for its flagship server products, such as SQL Server, Exchange Server, and SharePoint. Microsoft even left open the option that a support problem might need to be duplicated on physical hardware if support technicians couldn't determine the nature of the problem in a virtual environment. Adding to Microsoft's weak support story was the fact Virtual Server 2005 R2 was its virtualization product during the early days of Microsoft Office SharePoint Server (MOSS) 2007. Virtual Server 2005 R2 wasn't a hypervisor-based product and couldn't virtualize 64-bit guests, which limited supported environments to those running the 32-bit versions of MOSS 2007. This greatly curtailed the performance that could be achieved, particularly for the database role, which was the most resource-intensive and could take advantage of the 64-bit architecture the most. In addition, web front ends typically required significantly more memory than a 32-bit platform.

Two significant developments changed this story. The first was Microsoft's release of a 64-bit–capable hypervisor named Hyper-V. The second was the development of a program called the Server Virtualization Validation Program, which outlined Microsoft's official support stance on running its products on third-party hypervisor virtualization platforms. This program, outlined in the Microsoft article "Support policy for Microsoft software running in non-Microsoft hardware virtualization software" (support.microsoft.com/kb/897615), allowed for support of Microsoft products on third-party virtualization products that were validated by Microsoft and complied with certain criteria. These two developments opened the doors for Microsoft servers running on virtual machines (VMs) and gave peace of mind to organizations that needed to deploy supported virtualized solutions.

The 2007 wave of SharePoint products—which includes Windows SharePoint Services 3.0 (WSS 3.0) and MOSS 2007—was the first to gain broad virtualization support from Microsoft. However, in production, most organizations opt to virtualize only the web role and sometimes the query role. Other roles aren't typically implemented for various reasons. For example, the index role is often implemented only on physical hardware because of heavy processor and memory constraints and the limitation of one index server per Shared Services Provider.

Microsoft's official SharePoint 2010 support stance is that any SharePoint role or service is supported for hardware virtualization. SharePoint 2010 is positioned as a great version to virtualize because of the virtualization technology advances and reduction of disk I/O requirements for the indexing and search components. In addition, advances in hardware virtualization make it easier to virtualize I/O intensive applications such as SQL Server, allowing the SharePoint database role to be more easily virtualized. As a result, many organizations are looking at virtualizing their new SharePoint 2010 farms.

InstantDoc ID 125112

Microsoft

Efficiency you can feel, but can't touch.

Finally, a consolidated virtualized infrastructure, from the data center to the desktop, is within grasp. Well, not literally of course. After all, it's virtualized. Start with Windows Server® 2008 R2 with built-in Hyper-V™ and you can eliminate costly third-party software like VMware. Add SQL Server® 2008 Enterprise, with unlimited virtualization, and you just made it easier to eliminate racks of underutilized servers. Toss in System Center and you've centralized management across the enterprise all the way down to the application level. Translation? Flexible and dynamic virtualized infrastructures that help maximize ROI, reduce TCO and improve business continuity. Just don't let the efficiency go to your head.

To learn more about how server virtualization can make you more efficient, go to **itseverybodysbusiness.com/virtual**

Snap this tag to get the latest news on server virtualization or text VIRTUAL to 21710

Get the free app for your phone at **http://gettag.mobi**

Because it's everybody's business

of both Hyper-V and VMware hosts and guests through a single console.

Microsoft provides cost-effective virtualization licensing options for Windows Server, which lets organizations save significantly on Windows Server licenses when virtualizing servers. The three types of virtualization server licenses are:

- Windows Server Standard Edition, which allows a single physical OS environment (POSE) or a single virtual OS environment (VOSE) with each Standard Edition license. Note that a virtualization host that's dedicated to virtualization tasks doesn't consume a license, regardless if it's running Windows Server (such as in the case of Hyper-V).
- Windows Server Enterprise Edition, which allows for up to four VOSEs to be run at any one time on the host. Note that only running VMs are counted, so if a VM is shut down, it doesn't count against the four concurrent VOSEs permitted by the Enterprise Edition license.
- Windows Server Datacenter Edition, which is a per-processor license for the virtual host (e.g., a dual quad-core server would require two licenses) that grants you the right to run an unlimited number of VMs on the host.

These licensing options apply not only to Hyper-V but also to any hypervisor that's part of the Server Virtualization Validation Program. For organizations with a significant investment in virtualization infrastructure, buying the appropriate number of Datacenter Edition licenses to cover all the virtual hosts is the most cost effective.

Virtualization of the Web Role

Any SharePoint role or service can be virtualized, so which SharePoint server should you virtualize first? The best candidate is the SharePoint server that has the web role, which means it runs Microsoft IIS and handles all web requests sent to SharePoint. Table 1 shows resource guidelines for virtualized SharePoint servers that have the web and other roles.

As you can see in Table 1, a SharePoint server that holds only the web role (aka web server) should be allocated two virtual processors and a minimum of 6GB of RAM, along with a single VHD for the OS. If a web server needs to handle more web traffic, you can simply allocate additional web servers using the same specifications. The size of the host OS VHD should be at least 12GB plus the total amount of memory allocated to the VM, but it's good practice to size this volume larger (typically around 50GB to 100GB) to allow the host OS to grow in size.

Virtualization of the Application Role

The next likely candidates for virtualization are the SharePoint servers with the application role (aka application servers). Application servers can include various service applications, such as Access Service, PerformancePoint, and Managed Metadata. Generally, this excludes search services. Although they're technically service applications, they're typically classified under different server roles for architectural purposes. Thus, I'll talk about the search roles separately.

Any SharePoint role can be virtualized, so which SharePoint server should you virtualize first? The best candidate is the server with the web role.

As Table 1 shows, the typical virtualized application server consists of a VM with two virtual processors and a minimum of 6GB of RAM allocated to it. It needs a single VHD that's presized in the 50GB to 100GB range for the guest OS. Note that these numbers can vary, depending on how many service applications are installed on a single machine and how many people use the applications.

Smaller organizations sometimes combine the application and web roles on one SharePoint server. Combining the roles will increase the memory and processor requirements of the guest session.

Virtualization of the Search Roles

Third in line for virtualization are the SharePoint search servers, which are servers that hold one or more of the search roles. The search roles include the query, index, and crawl roles. The query role provides querying functionality, whereas the index role provides indexing functionality. SharePoint 2010 doesn't have the same single-index restrictions that Microsoft Office SharePoint Server (MOSS) 2007 has, which makes this role more scalable and allows for more distributed deployment models. The crawl role represents the crawler component used by SharePoint to crawl documents for search purposes. Multiple crawl components can be created on different servers for redundancy.

A typical virtualized search server consists of a VM with four virtual processors and 8GB of RAM allocated to it (see Table 1), assuming that SharePoint 2010's out-of-box search functionality is being used. If FAST Search Server 2010 is being used, the RAM requirements will be in the 12GB to 16GB range. Like the application server numbers, the search server numbers can vary, depending on how many items are being indexed and how heavy the search requirements are.

The search server needs a single VHD that's presized in the 50GB to 100GB range for the guest OS and another VHD for the index and query corpus. The size of this VHD will vary, depending on how much full text is being indexed from various sources.

Smaller organizations sometimes use one SharePoint server for the search and web

Table 1: Resource Guidelines for the SharePoint Server Roles

Roles	Virtual Processors	Minimum RAM
Web role only	2	6GB
Application role only	2	6GB
Search roles only	4	8GB
Combined web, application, and search roles	4	10GB
Database role	4	8GB

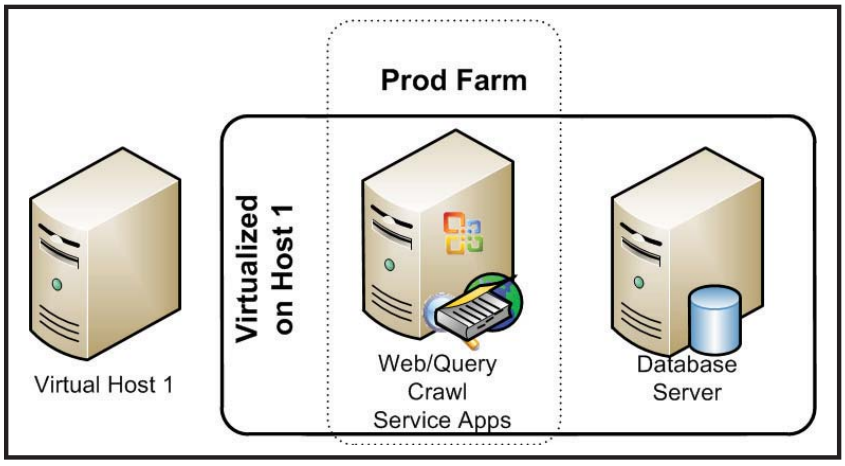

Figure 1: Small virtualized SharePoint 2010 environment with one virtual host

roles. Combining these roles can increase the memory and processor requirements of the guest session.

Combining All the Roles on a Virtualized Server

Many organizations combine the web, application, and search roles on a single virtualized SharePoint server. This is often the case in small organizations that want to deploy SharePoint across two guest sessions to be highly available, but have a small number of guests.

Although combining these roles results in additional load on an individual server session, many of the same processor and memory guidelines that apply to a dedicated web server apply to a combined server. As Table 1 shows, a typical virtualized combined server consists of a VM with four virtual processors and 10GB to 16GB of RAM allocated to it, depending on how many users the system will support. It has a single VHD presized in the 50GB to 100GB range for the guest OS and another VHD for the index and query corpus.

SharePoint administrators familiar with MOSS 2007 might be dismayed at the memory requirements of SharePoint 2010, but the fact is that SharePoint 2010 requires much more memory than previous versions. RAM requirements can be lessened, however, by turning off service applications that aren't required by the business.

> The server with the database role is the last and most challenging to virtualize. It needs the lion's share of RAM and processor allocation.

Virtualization of the Database Role

The SQL Server database role is the last and most challenging server role to virtualize. The server with the database role (aka database server) needs the lion's share of RAM and processor allocation. A minimum of four virtual processors and 8GB of RAM should be allocated to the database server. For best performance, though, at least 12GB of RAM should be allocated.

Like SharePoint VMs, SQL Server VMs require either fixed-sized or pass-through VHDs. The same disk considerations that would apply to physical SQL Server machines apply to virtual SQL Server machines, so be sure to allocate enough disk spindles for the database and log volumes. In addition, be sure to follow standard best practices for SharePoint–SQL Server optimization, such as presizing tempdb and moving it to fast disk volumes.

Keep in mind that these guidelines are simply guidelines. Actual performance will be dictated by the type of disk, hardware architecture, and other factors. Some organizations calculate their hardware requirements, then simply add RAM or reduce the number of databases on a single SQL Server session.

Microsoft supports both database mirroring and clustering as high-availability options in a virtualized SQL Server environment. In addition, host failover options such as Hyper-V Live Migration are supported for SQL Server VMs. One fact to note, however, is that all SQL Server databases within a SharePoint farm need

Table 2: Deployment Specifications for a Small Virtualized SharePoint Environment

Server	Memory	Processors	Disk
Virtual host	24GB RAM	2 quad-core (8 cores)	C drive: OS; Windows Server 2008 R2 with Hyper-V; dedicated volume (50GB) D drive: dedicated volume for OS VHDs E drive: dedicated volume (500GB) for SQL Server database VHDs F drive: dedicated volume (100GB) for SQL Server log VHDs
SQL Server server	10GB RAM	4 virtual processors	C drive: OS; fixed-size VHD (100GB) D drive: fixed-size VHD (100GB) for the SQL Server logs E drive: fixed-size VHD (500GB) for the SQL Server data
SharePoint web/ query/app server	8GB RAM	4 virtual processors	C drive: OS and transport queue logs; fixed-size VHD (100GB) E drive: fixed-size VHD (100GB) for indexing and querying

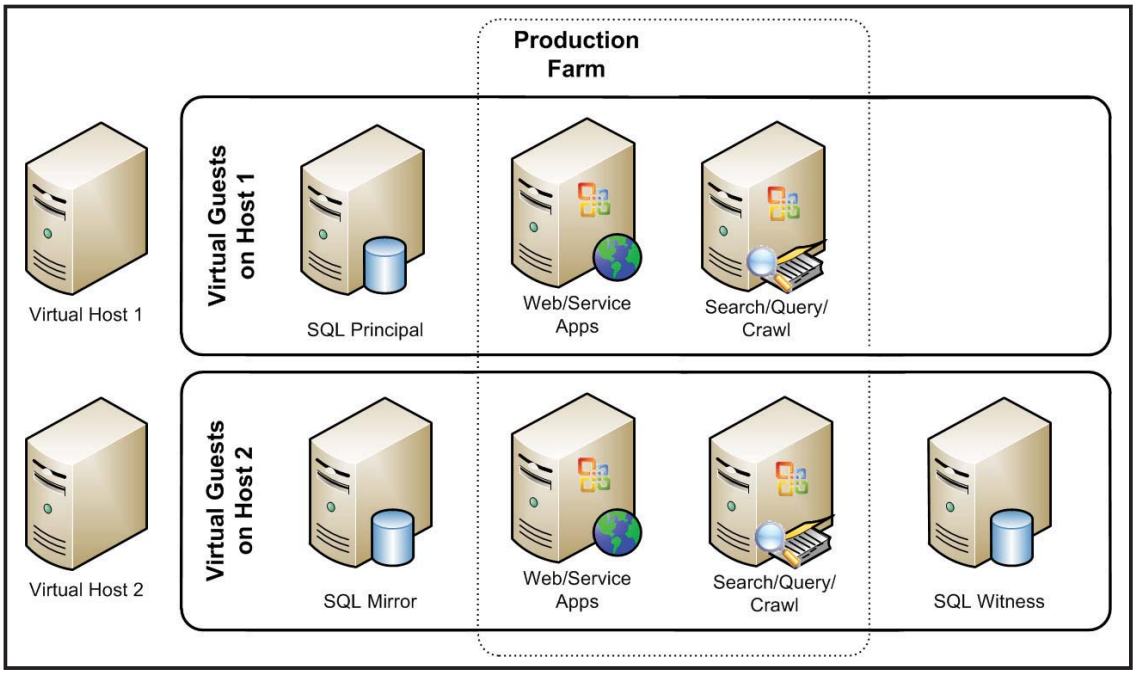

Figure 2: Mid-sized virtualized SharePoint 2010 environment with two virtual hosts

to be restored from the same point in time as the other databases. This applies to virtualization snapshot technology or SAN-based snapshots of SQL Server databases.

Sample Virtualized SharePoint 2010 Architectures

There are many ways to deploy SharePoint 2010 in a virtualized environment. However, some designs are more common than others and reflect common needs across many organizations. For example, high availability is becoming a must for the crucial document management and collaboration functionality in SharePoint. All the new high-availability options in SharePoint 2010 are available for virtual environments and can be easier to deploy because of the flexibility that virtualization provides.

Figure 1 illustrates a small virtualized SharePoint 2010 environment with all components running on a single virtual host. This type of deployment doesn't have any built-in high availability or disaster recovery, but it's the simplest environment to set up and it can still take advantage of virtualization benefits and scalability. Table 2 shows sample server specifications for an environment of this size. These specifications assume there are 500 active users in the environment.

Figure 2 illustrates a virtualization architecture that provides a high level of availability, disaster tolerance, and scalability for an environment with 2,000 active users. The entire SharePoint environment is deployed across two virtual hosts, which provides for high availability of the environment. SQL Server databases are mirrored from one virtual guest to another, and a witness SQL Server instance monitors the principal SQL Server instance, providing for automatic failover in the event the virtual host or virtual guest fails.

These high-availability and disaster-recovery options are possible without the need for shared-storage, SAN, or host-availability solutions. Table 3 lists the

Table 3: Deployment Specifications for a Mid-Sized Virtualized SharePoint Environment

Server	Memory	Processor	Disk
Virtual hosts	48GB RAM	2 quad-core (8 cores)	C drive: OS; Windows Server 2008 R2 with Hyper-V; dedicated LUN (50GB) D drive: dedicated LUN for VHDs Raw volume: dedicated LUN (100GB) for SQL Server logs Raw volume: dedicated LUN (2TB) for SQL Server databases
SQL Server servers	16GB RAM	4 virtual processors	C drive: OS; fixed-size VHD (50GB) D drive: pass-through dedicated LUN (100GB) for SQL Server logs E drive: pass-through dedicated LUN (2TB) for SQL Server data
SharePoint web/ app servers	12GB RAM	2 virtual processors	C drive: OS; fixed-size VHD (100GB)
SharePoint search/ query servers	12GB RAM	2 virtual processors	C drive: OS; fixed-size VHD (100GB) D drive: fixed-size VHD (200GB) for indexing and querying
SQL Server witness	2GB RAM	1 virtual processor	C drive: OS; fixed-size VHD (50GB)

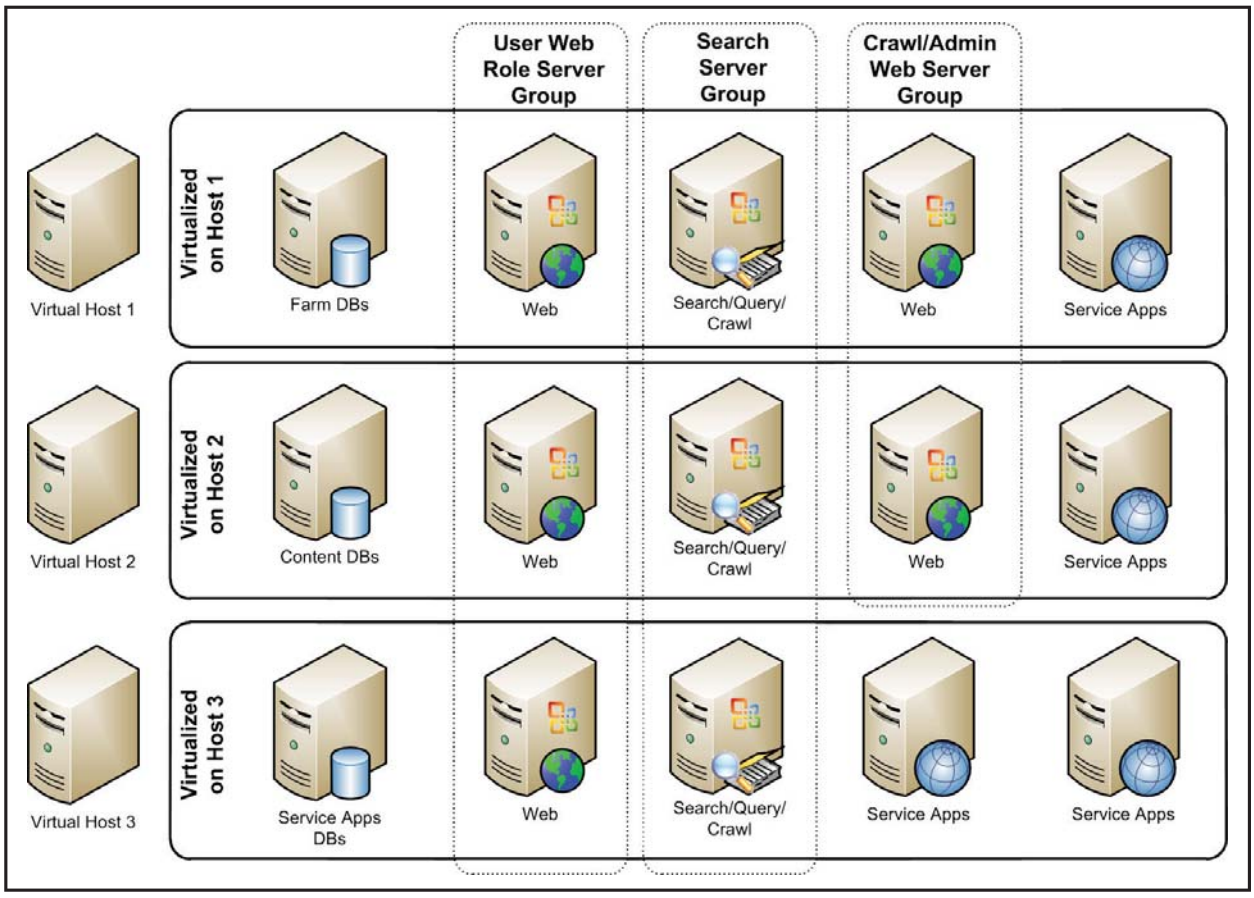

Figure 3: Large virtualized SharePoint 2010 environment with three virtual hosts

sample virtual host and guest architecture guidelines for the solution in Figure 2.

Virtualization technologies allow for a high degree of scalability and aren't limited to small and mid-sized organizations. For example, the architecture shown in Figure3 allows for tens of thousands of SharePoint users, full disaster tolerance, and high availability, all with the high performance expected from SharePoint. In this particular model, multiple SQL Server machines are used for the various SharePoint databases, with one used for content databases, one for service application databases, and one for the search databases. Server groups are created for different SharePoint server roles, and the web tier is broken into two components: one for users and another for crawl and administration. In this example, host-based failover solutions such as Hyper-V Live Migration could also conceivably provide for failover of individual guest sessions between failed hosts. The virtual host and guest architecture guidelines for the solution in Figure 3 would be similar to the guidelines in Table 3, but scaled larger.

With the proper virtualization architecture, you can deploy a fault-tolerant and high-performance SharePoint 2010 environment.

These three samples architectures illustrate some of the potential design options that are available for a virtual SharePoint environment. Every environment is unique, and specifics will vary based on business and technology needs. However, you can use these sample architectures as a starting point for developing a high-performance virtualized SharePoint 2010 environment.

The Virtualization Advantage

Server virtualization can provide significant advantages and can let SharePoint architects design highly available and disaster-tolerant environments more easily than could be done solely on physical hardware. In addition, virtualized environments have consolidation, optimization, and cost-saving benefits that make them ideal for many organizations. With proper thought into host and guest virtualization architecture, you can deploy a fault-tolerant and high-performance SharePoint environment that lets you fully capture the benefits of virtualization for your organization.

InstantDoc ID 125111

Citrix XenApp Management Pack for SCOM

HERMES SoftLab announced the imminent release of a Citrix XenApp Management Pack for Microsoft System Center Operations Manager (SCOM). The new XenApp Management Pack complements the XenServer Management Pack, adding support for Citrix XenApp. Both management packs were developed in cooperation with Microsoft and Citrix to provide comprehensive health, performance, and topology information, all natively integrated into SCOM. The packs also offer out-of-box monitoring knowledge, tasks, and reports created specifically to enable day-to-day supervision. To learn more, visit www.hermes-softlab.com.

Sans Digital Offers Potential 224TB Storage NAS

Sans Digital has announced updates to its **EliteNAS EN208L+BXE, EN212L+BXE,** and **EN316L+BXE** storage devices. The updated models support both SAS and SATA hard drives and can be expanded by connecting up to seven units together. With 16 bays per device and 7 devices, this provides the potential to link 112 hard drives, totaling 224TB if you assume 2TB of storage per hard drive. The EliteNAS series also supports NAS for data sharing and iSCSI for virtualization applications. Optional 10Gb Ethernet network interface controllers are available. The EliteNAS EN316L+BXE (which offers 16 bays) costs $5,450. To learn more, visit www.sansdigital.com.

PRODUCT SPOTLIGHT

Dot Hill Offers Disaster Recovery for Small Businesses

Dot Hill Systems expanded its channel program with the introduction of AssuredSAN disk-to-disk data protection appliances. The Dot Hill **AssuredSAN 3000 Series** features a suite of disaster recovery applications, including Dot Hill's new AssuredRemote remote replication software. With volume copy, snapshot, and remote replication, this solution makes affordable, comprehensive disaster recovery solutions available to small and medium-sized businesses.

AssuredRemote offloads backup operations from critical application servers, creating little to no impact to production servers by asynchronously copying and synchronizing data for up to 16 volumes simultaneously, directly between two AssuredSAN 3000R appliances via the local storage area network (SAN) or wide area networks (WANs). Replication can operate via the Fibre Channel ports and/or the iSCSI ports in the case of the dual protocol AssuredSAN 3900 models, and supports high-availability failover using dual controllers.

"The response to our AssuredSAN launch in September was extremely positive," said David Zimmer, vice president of worldwide channel sales and marketing, Dot Hill. "Customers told us that they would like to see more data protection features such as AssuredSnap and AssuredCopy in future-generation products. The AssuredSAN 3000 Series is the result; by adding AssuredRemote to our new 8Gb Fibre Channel-based 3000 Series architecture, we're providing SMB customers with affordable, midrange-featured data protection appliances."

The Dot Hill AssuredSAN 3000 Series is available through Dot Hill channel partners. List prices are expected to start at $20,200 for a dual controller Fibre Channel system with 500GB SATA drives. To learn more, visit www.dothill.com/3000series.

Siemon Announces Security Features for Z-MAX Cabling Outlets

Siemon has announced a few security enhancements for its **Z-MAX** network cabling outlets. First, the connectors are now available with an optional spring-loaded hinged door, which protects the outlet's internal components and connector mating surfaces from exposure to environmental contaminants such as dust. Also, using Z-TOOL, Z-MAX door-equipped outlets can be terminated in 60 seconds or less from start to finish, including cable preparation. Siemon's Z-MAX line includes category 6A/class EA systems in both shielded and unshielded configurations as well as category 6/class E UTP. The full Z-MAX 6A channel consists of Z-MAX 6A outlets and patch panels, Siemon category 6A cable and patent pending Z-MAX 6A patch cords that feature a precisely tuned printed circuit board (PCB) within every plug for enhanced signal performance. To learn more, visit www.siemon.com/us/zmax.

Advanced Emulation for Cyber Security Capability Development

Scalable Network Technologies announced the introduction of **EXata/cyber,** a new software tool designed to support and accelerate development of cyber security capability for communication networks. EXata/cyber provides a robust emulation

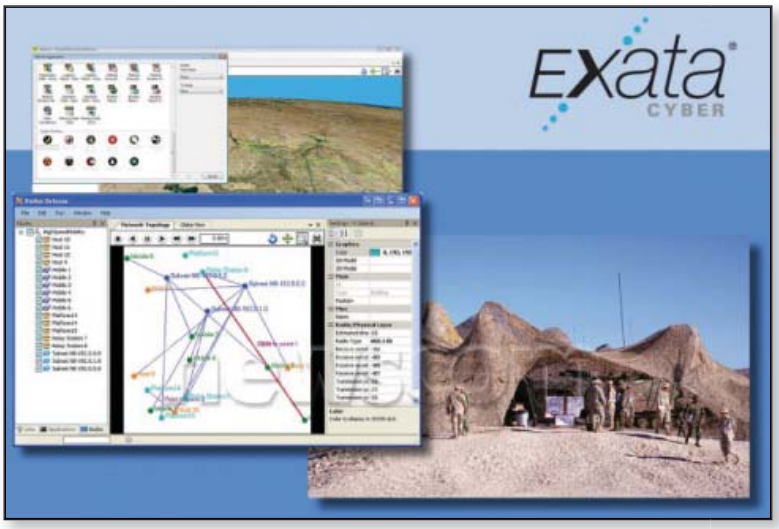

platform that can expose vulnerabilities that threaten communication networks while letting you safely test and develop countermeasures. EXata/cyber makes it possible to vet networks through creation of a "software virtual network" (SVN). SVNs are exact digital replicas of physical networks in virtual space—indistinguishable from a real network. EXata/cyber comes in two parts: the main EXata emulation engine that creates a digital replica of the user's target network, and the Connection Manager that runs on their operational systems. To learn more, visit www.scalable-networks.com.

Lock Your USB Ports from Malicious Use

AC Element Company has released **MyUSBOnly 6.0,** a Windows USB control and security application that prevents data theft through computer USB ports. With MyUSBOnly, you create a whitelist of USB devices that you want to allow. The product also lets you review the event log, examine recent activity, view statistics, or search for specific, detailed security information. The latest version of MyUSBOnly includes the ability to deploy and configure MyUSBOnly remotely on all the computers in the organization, reducing administrative workload. You can also configure the software to send email alerts of all USB access activity. MyUSBOnly costs $29.90 for a single-user license. To learn more, visit www.myusbonly.com.

Recover Office Passwords Faster

Accent has released **OFFICE Password Recovery 3.5.** The latest version uses the latest graphics cards to recover lost Microsoft Office 2007 and 2010 passwords quickly. According to the vendor, Accent OFFICE Password Recovery 3.5 is the only tool of its kind that supports both NVIDIA and ATI graphics cards. By linking all a computer's existing graphics cards into one unit, Accent OFFICE Password Recovery gives users top speeds for recovering Microsoft Office 2007/2010 passwords. Accent OFFICE Password Recovery runs on all versions of Windows from XP and on. Pricing starts at $60 for a Basic license. To learn more, visit www.accentsoft.com.

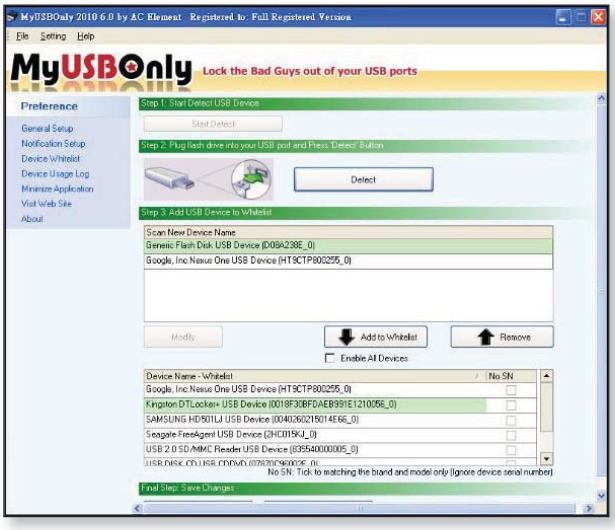

Paul's Picks

www.winsupersite.com

SUMMARIES of in-depth product reviews on Paul Thurrott's SuperSite for Windows

Windows Phone

PROS: Innovative, unique take on the smartphone; killer hardware specs

CONS: Unclear whether customers will embrace another new platform

RATING: ◆◆◆◆◇

RECOMMENDATION: When Microsoft revealed that it was going to dump Windows Mobile and start over, I was surprised. When it revealed that it was going to build off the underappreciated Zune platform, I was pleased. But when it revealed that it was tossing out the tired, apps-based model used by the iPhone and its many copiers and making a truly innovative user experience, I was floored. Whether this new platform will be embraced by users or developers remains to be seen. But the work looks solid, and even in these early days, it appears that Microsoft is on to something: Developers have a mature .NET and Silverlight-based environment to work in. And users will have a touch-tastic new UI that might eventually even be seen on non-phone devices. Color me excited: Windows Phone is a platform to watch.

CONTACT: Microsoft • www.microsoft.com

DISCUSSION: www.winsupersite.com/mobile/wp7_preview.asp

iPhone OS 4.0

PROS: Multitasking ability; better enterprise support; more customization

CONS: The above features should have been in there all along; multitasking isn't complete

RATING: ◆◆◆◆◇

RECOMMENDATION: With its iPhone OS 4.0 update, due in mid-2010, Apple does something new with its smartphone platform: It's aping the competition. Whether this is a sign of maturity or of desperation is unclear, but the iPhone is picking up a lot of familiar functionality. It's getting multitasking (like Windows Mobile/Windows Phone and Android); customization features (Android); support for multiple Exchange accounts (Android) and an integrated Inbox (WebOS); and enterprise Exchange features (RIM Blackberry, Windows Mobile). And it's even getting a Game Center, similar to the Xbox Live support on Windows Phone. But regardless of the sources of its functional inspirations, Apple's iPhone OS should continue its reign atop the smartphone market. With this update, a good thing is getting better.

CONTACT: Apple • www.apple.com

DISCUSSION: community.winsupersite.com/blogs/paul/archive/2010/04/08/iphone-os-4-0.aspx

InstantDoc ID 125031

REVIEW

PRTG Traffic Grapher

Having used only ISP-provided bandwidth monitors, I'd never set one up and didn't know how bandwidth monitors actually gather the data they use to display traffic graphs and statistics. Unfortunately, the bandwidth monitor that my company's ISP uses shows information only for the Internet connection at our main location; we recently needed data for all 21 locations before committing to a private network upgrade. A consultant friend recommended Paessler's **PRTG Traffic Grapher.**

This valuable tool supports SNMP monitoring via SNMPv1, SNMPv2, and SNMPv3; Cisco Systems' NetFlow protocol, packet sniffing (which provides data about applications and devices), or latency (which uses ICMP to identify overloaded devices or network segments). Paessler's recommendation for PRTG Traffic Grapher is to use SNMP in most cases, or NetFlow for large or traffic-heavy Cisco networks. To use the packet sniffing sensor, you must plug into a switch's monitoring port or install PRTG Traffic Grapher on a computer acting as a router between the subnets you want to monitor. In addition, Windows Server 2008 doesn't support the packet sniffing sensor.

PRTG Traffic Grapher's system requirements are pretty loose. According to the company's website, Windows XP or later, either 64-bit or 32-bit, will work. Hardware-wise, you need a server, PC, or virtual machine (VM) with the equivalent performance of an average computer built in 2007. To view the PRTG Traffic Grapher web console, you need Internet Explorer 8.0, Google Chrome, Mozilla Firefox, or Apple's Safari. I installed the product on Windows Server 2008, Windows Vista, and XP without any problems.

To install PRTG Traffic Grapher, download the executable file from Paessler's website and run it on a Windows client. Along with the desktop application, the software installs a web-based management interface. If you're already running a website on the machine you want to install the software on, you can change the port for the PRTG web server in the program options.

Next, you need to decide what interfaces you want to monitor. I wanted to know the amount of bandwidth being used on the WAN and by which sites. We use SonicWall firewalls at each of our 21 locations to form our IPsec VPN. Because our SonicWall devices can all be managed with SNMP, and all site traffic passes though their interfaces, monitoring them would give me the information I needed. I logged on to each device and configured SNMP, pointing to PRTG Traffic Grapher.

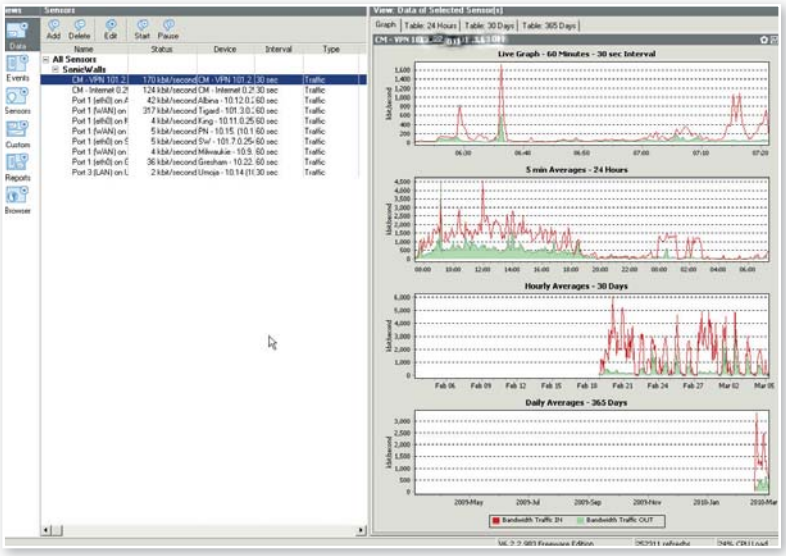

Figure 1: PRTG Traffic Grapher console

Once the endpoint devices are ready to share their SNMP data, you need to configure PRTG Traffic Grapher to collect the data. From the Start menu, select PRTG Traffic Grapher to launch the console application, which Figure 1 shows. In the console's Data section, select Edit, Add Sensor. Click through the wizard and enter your endpoint device's IP address when prompted. If PRTG Traffic Grapher can't communicate with a device via SNMP, the software will prevent you from completing the Add Sensor section. I ran into this problem a couple of times when I forgot to enable SNMP for the SonicWall LAN interface that I wanted to monitor.

Next, you can customize your configuration. You can configure error notification via email, set a limit on your bandwidth graphs to remind you of a particular length's maximum, set up reports to automatically email at designated times, and so on.

Since installing PRTG Traffic Grapher, I've used the program to perform several useful tasks. I gathered all the data necessary to determine the size of circuits we'd need at each location for the network upgrade, identified a device that was being overloaded with network traffic and needed to be replaced, and tracked down a user who decided to plug his personal laptop into the corporate network and start hogging bandwidth.

Even if you have network monitoring in place, you might want to take advantage of PRTG Traffic Grapher's ability to generate real-time and historical data about traffic usage. The software is easy to set up and use, and the product provides valuable data about networks of any size.

InstantDoc ID 125064

PRTG Traffic Grapher

PROS: Minimal installation requirements; easy to install and use; alerts and reports included; inexpensive

CONS: Can't alert by modem in the event of network outage

RATING: ◆◆◆◇◇

PRICE: 10 nodes, free; 100 nodes, $380; 500 nodes, $995; 1,000 nodes, $1,595; unlimited nodes, contact vendor

RECOMMENDATION: I recommend the product for anyone working on a network with more than one location or in an environment with a lot of network traffic.

CONTACT: Paessler • +49-911-7872497 • www.paessler.com

Nate McAlmond | mcalmond@gmail.com

Group Policy Change Reporter

Have you ever had the following conversation with your boss? "The Group Policy accident last week cost us quite a few hours of production. We need to know who pulled the trigger, and when." Fortunately for you, Windows Server, like most modern network OSs, has robust logging built in. Unfortunately for you, your window into this logging comes via the Event Log tool. Using this tool to track down events that happened today, let alone last week, is nearly impossible. If this scenario sounds familiar, you need a tool like NetWrix's **Group Policy Change Reporter.**

Group Policy Change Reporter is one program in a family of useful tools. NetWrix offers many administrative and reporting tools, such as Exchange Change Reporter, File Server Change Reporter, Disk Space Monitor, and many others.

I've installed several NetWrix applications in the past year, and I've been quite impressed. The process is always simple and intuitive. To install Group Policy Change Reporter, I had to first install the .NET Framework 2.0 and Group Policy Management Console (GPMC). After I completed these steps, the actual software installation took only a few minutes. You can install the application directly on the domain controller for smaller domains; for large domains, you might want to use a dedicated utility server (e.g., a server that hosts your antivirus and Windows Server Update Services—WSUS—updates, etc.).

When the installation is finished, a configuration page displays options such as the domain you want to monitor, a location for the data, the amount of time you want to keep the logs, and an email account the reports should be sent to.

To test the tool, I created a simple organizational unit (OU) structure and proceeded to implement some new Group Policy Objects (GPOs). First, I created a GPO called PC – XP – Wait for Network (which is an important Group Policy setting if you want to deploy software to computers via GPO). After this Group Policy setting was implemented, I unlinked the Group Policy setting from the OU to simulate a junior administrator making a change to a production network.

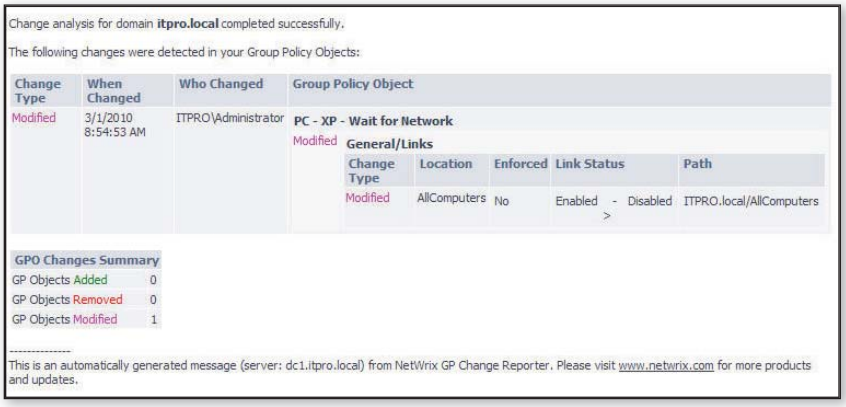

Change analysis for domain **itpro.local** completed successfully.

The following changes were detected in your Group Policy Objects:

Change Type	When Changed	Who Changed	Group Policy Object
Modified	3/1/2010 8:54:53 AM	ITPRO\Administrator	PC - XP - Wait for Network

Modified — General/Links

Change Type	Location	Enforced	Link Status	Path
Modified	AllComputers	No	Enabled -> Disabled	ITPRO.local/AllComputers

GPO Changes Summary	
GP Objects Added	0
GP Objects Removed	0
GP Objects Modified	1

This is an automatically generated message (server: dc1.itpro.local) from NetWrix GP Change Reporter. Please visit www.netwrix.com for more products and updates.

Figure 1: Ad-hoc Group Policy Change Reporter report

To view a report of my changes, I could have waited until the daily 3:00 A.M. report ran. But to speed things up, I decided to run the report manually via Scheduled Tasks. (On the Start menu, select All Programs, Accessories, System Tools, Scheduled Tasks.) I then used the NetWrix Enterprise Management Console to run an ad-hoc report. After a few seconds, the report in Figure 1 opened in Internet Explorer. As you can see, the report clearly shows that I modified the Group Policy setting, changing the Link Status from Enabled to Disabled. It also shows that the user who performed the change is "Administrator," which brings up a great point about accountability: Never let your administrators log on as "Administrator"—if you do, you'll never really know who performed a specific task. You can have your reports emailed to you every morning, which will give you a nice 24-hour snapshot of what your administrators have been up to.

Setting up Group Policy Change Reporter's reporting structure is extremely simple. If you require the product's Advanced Reports, you'll need to configure a SQL Server machine with SQL Server Reporting Services (SSRS). Fortunately, a configuration page walks you through the entire process. You can select from 13 built-in reports, such as Account Lockout Policy Changes and Security Policy Changes.

Support for NetWrix products is available through a free support forum, a robust and searchable Support Knowledge Base, a ticketing system for online support, or by toll-free phone (United States only). If you need help setting up a large implementation, NetWrix offers contract services that can ease deployment, offer customization, and provide training.

Group Policy Change Reporter is a great little program, with sister applications that augment its functionality (e.g., Active Directory Change Reporter, Exchange Change Reporter). The product is inexpensive and simple to set up. If you need easy, hassle-free reporting to keep track of what's happening on your network, give Group Policy Change Reporter a try.

InstantDoc ID 125023

Group Policy Change Reporter

PROS: Extremely easy setup for simple reporting; easy-to-navigate interface via an MMC snap-in

CONS: None found

RATING: ◆◆◆◆◆

PRICE: $2.50 to $4.00 per enabled user, depending on number of users; discounts available

RECOMMENDATION: If you're tired of always wondering who changed a Group Policy setting and you need an extensive reporting tool, then Group Policy Change Reporter is the solution to use.

CONTACT: NetWrix • 888-638-9749 • www.netwrix.com

Eric B. Rux | ebrux@whshelp.com

REVIEW

InterMapper 5.2

One of the simplest physics concepts to swallow is that of *entropy.* Entropy is a level of randomness, or disorder, within a system. Taking ample poetic license, network administrators can contextually appreciate the law that entropy always increases in an active system. Admins see this process in action every day as we carry out repairs and maintenance on running infrastructures. In addition, a corollary of the law of entropy is that a system or device will manage to break itself if you don't break it first. Therefore, it can't hurt to plan for this eventuality. Dartware helps restore some order to the IT universe with its **InterMapper** network monitoring software.

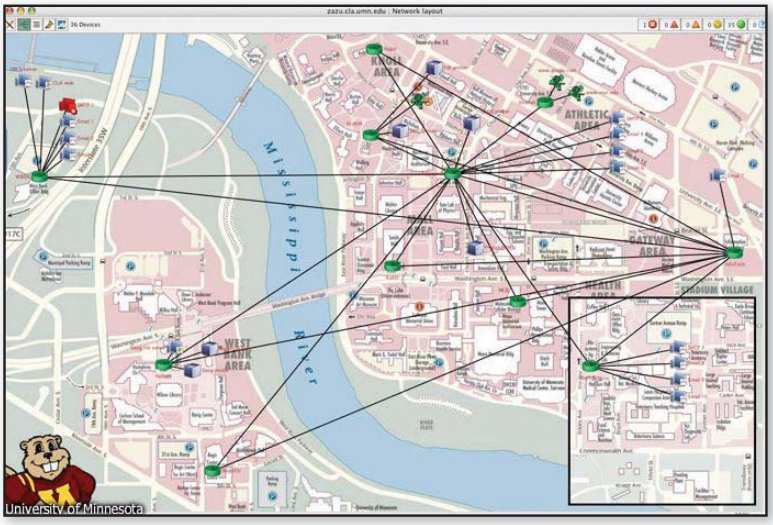

Figure 1: Using an actual map as an InterMapper background

InterMapper's initial installation is brief and uneventful, only prompting the user for a license key (if you forego the demo and jump right in). Network gurus from all backgrounds will be happy to know that the product includes support for most platforms, including Windows 2000 Server or later, OS X, Solaris, and several major Linux distributions, including Red Hat, SUSE, and Debian. InterMapper is Java-based, which will help ensure that this ubiquity perseveres. Installation requires a Java Virtual Machine (JVM), which is freely and widely available. Other minor requirements include 512MB of RAM, 50MB of disk space for the program, and 1GB of space for historical data storage and analysis.

After you complete the licensing steps, InterMapper fires up with an initial demo network map that uses several of the product's more interesting features. (Maps are the primary vehicle InterMapper uses to convey information.) InterMapper's maps can provide visual context through the use of various backgrounds, whether a scanned diagram drawn on a paper napkin or a satellite image downloaded from Google Maps. These backgrounds make it easy to move your network icons around, which provides additional value if you use an actual map for the background (as in Figure 1) and plot devices by their positional coordinates.

Excellent icon sets, flashing colors, and device info displays are standard network monitoring software features—Dartware goes further, with manual arrangement of icons to help you visualize grouping and hierarchy, as well as numerous options for forced arrangement by topology (e.g., star, bus, ring, and cell—or "organic"—layouts). These features can be applied to subsets of devices to allow for the desired combination of techniques. In addition, Dartware maintains an excellent gallery of screenshots from which users can get inspiration for their own maps.

InterMapper Flows is a purchased add-on that takes InterMapper's visualization aspect to another level by analyzing NetFlow, sFlow, and jFlow data aggregated from higher-end routers and managed switches. This data gives you a statistical view of your network traffic, akin to ntop.

The product does have some minor flaws. For example, it lacks a unified interface for starting or stopping network scans; some scans (or "reprobes") can occur with little or no indication when they start or stop. In addition, InterMapper's scan display could be more detailed—for example, sometimes the product begins a scan in the link-local address range due to an unconfigured interface, which can cause an exhaustive scan if you aren't paying attention.

InterMapper has the basic displays that other packages include, but the product adds visual features and centers around them. The software also works on a wide variety of platforms and has a small arsenal of sensors (including WMI-based probes in InterMapper 5.2). The InterMapper Flows add-on extends the product's range of application by adding NetFlow, sFlow, and jFlow compatibility. InterMapper stands alone in its unique, engaging presentation of the data we've all come to know and loathe.

InstantDoc ID 125040

InterMapper 5.2

PROS: Vast multitudes of probe types; visual interface is the primary display, rather than a dashboard panel; great SNMP/NetFlow/sFlow/jFlow support

CONS: Low-end pricing is fairly high per device; scanning operations would benefit from more detail

RATING: ◆◆◆◆◆

PRICE: $500 for 25 devices (doesn't include InterMapper Flows); $3400 for 500 devices; additional packages available

RECOMMENDATION: Although small operations can get by with more basic monitoring, InterMapper's features are well suited for medium to large organizations, especially those spanning considerable geographical area.

CONTACT: Dartware • 877-276-6903 • www.intermapper.com

Brandon Carse | bcbigb@gmail.com

3X Systems Remote Backup Appliance

The **3X Systems Remote Backup Appliance** (RBA) is a backup workhorse that lets you store—and restore—multiple versions of backed-up files. In addition, the RBA lets you configure exclusions so that users don't, for example, fill up backup storage space with personal MP3 files or videos. The device allows for quota management and supports all the major Windows clients in use today. As an additional benefit—for small businesses and SOHO installations—the RBA can perform online backups of Exchange Server storage groups and SQL Server databases, making it a complete backup system for the average SMB.

I tested the 500 Series RBA, which offers up to 500GB of RAID1 storage, a 1Gb Ethernet NIC, availability in either a 10.5" × 13" × 8" portable cube or a 1U rack-mountable chassis, and recommended support for as many as 50 users per unit and as many as five units. Upon opening the package, I was happy to find an included USB flash drive containing everything I needed to configure the RBA. After connecting the RBA to Ethernet and power, I walked over to my desktop computer to begin the configuration. I inserted the USB flash drive in my Windows 7 computer, and the admintool .exe file ran automatically.

I was presented with a firewall notification for javaw.exe and approved communications for the executable. After the system detected the RBA, I right-clicked it and selected Launch Manager, as the included *Quick Start Guide* instructed me to do. Because the IP address wasn't properly configured on the RBA, the connection didn't succeed. (I don't use DHCP on my test network.) I followed the instructions included for configuring a static IP address and was able to connect to the device easily to perform the initial configuration.

The initial configuration wizard is browser-based and simple to follow. It asks you to configure several items, a system name for the appliance (I used "RBA1"), the admin account password, and primary and secondary contact information. You also need to configure the time zone and network settings. 3X Systems provides a complimentary SMTP server for notifications so that you don't have to configure the SMTP server if you don't want to use your internal services.

Once the initial configuration is complete, you can configure the backup agent on the client computers. The backup agent performs scheduled backups to the RBA nightly or even throughout the day. The scheduling engine is flexible in this aspect. You can perform several functions from the administrative interface's start page, including appliance administration and client provisioning.

More complex features take this device from a consumer NAS solution to a business-class backup solution. The proprietary deduplication algorithm does an excellent job of ensuring that only one copy of a unique file is placed on the device. You can adjust the utilized network ports to work with your existing network environment and services. Also, you can maintain multiple versions of backed-up files that users or support personnel can restore.

While performing backups, I noticed some intermittent lags in performance. Overall, however, the machine boasts NAS-level performance while implementing the aforementioned deduplication, encryption, and file versioning. Impressive! I backed up more than 1.2GB of data, and the process completed in less than 15 minutes across a busy wired LAN. The second backup of the same data set took less than 60 seconds after I intentionally changed about 50MB of data to represent a few days' work.

Be sure to configure each RBA on a management port that's separate from any other device. (The default is port 443 for SSL.) If you're already using SSL, you'll need to change the TCP port number in the RBA and possibly in your port forwarders at the network perimeter. With the proper configuration, you can push restores to remote clients across the Internet so that the user needn't master the file-restoration process. This push restore feature is my favorite feature in the entire system.

Although the RBA is effectively a computer with a standard gigabit Ethernet connection, it's so much more than a simple computer with basic backup software installed. The RBA is a true backup appliance that should provide sufficient backup solutions for SMBs or departments requiring centralized backup of client computers. The price-to-feature comparison results in a good investment for your data backup and recovery needs. (The company also provides a Tera Series and an Enterprise Series—the more powerful series models provide for more storage space, more computing power, and greater network bandwidth capabilities.)

InstantDoc ID 125182

3X Systems Remote Backup Appliance

PROS: Fast, simple configuration; excellent support for encryption and compression; proprietary deduplication features significantly reduce the required storage space

CONS: Configuring remote access can be a bit of a chore if you have existing SSL solutions on your network

RATING: ◆◆◆◆◆

PRICE: $2,495

RECOMMENDATION: 3X Systems' 500-series RBA device—a centralized backup solution that backs up clients and servers without requiring continual intervention from a support specialist—is an excellent solution for the SMB.

CONTACT: 3X Systems • 866-478-3131 • www.3x.com

Tom Carpenter | (carpenter@sysedco.com)

SharePointPro CONNECTIONS

Your SharePoint expert community

SharePointPro Connections

provides guidance to help align the technology with business requirements. You'll get real-world advice from professionals and peers who share their experience administering and developing in SharePoint.

SharePointPro Connections is the independent voice on SharePoint technology; expert authors provide readers with the field-tested information they need to enable content and image management, collaboration, and workflow solutions tailored to their business needs.

Subscribe FREE to the only magazine dedicated to all things SharePoint.

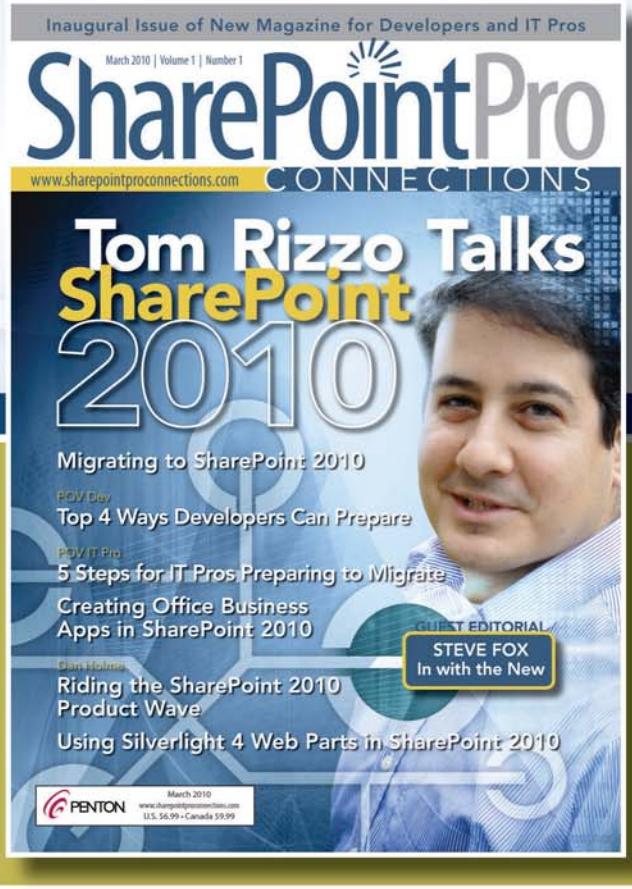

Upcoming articles include:

- Migrating to SharePoint 2010
- Using SharePoint with Visual Studio
- Getting more out of SharePoint

ORDER YOUR FREE SUBSCRIPTION TODAY!

sharepointproconnections.com/go/subscribetoday

Windows IT Pro

Security Special Report

by John Savill

Security is at the forefront (no pun intended) of nearly every IT manager's mind today. As more organizations digitally store information and increasingly rely on their IT infrastructures, they must ensure that their resources are protected from malicious parties while they maintain an intuitive and transparent end-user experience.

One guiding principle you will hear time again with security is "defense in depth"—the idea that multiple layers of defense provide the best protection. Then if one layer has a weakness, the next layer still offers security. Just as cars have frames to absorb impact, and then air bags to give extra protection, your IT resources need layers of security that complement each other and add redundancy to your overall security approach.

In this guide to system security, we examine three different security areas: 1) primary Microsoft security solutions that make up the Forefront family, 2) how best to protect your virtual infrastructure, and 3) how to offer secure services on the Internet.

SYMANTEC IS

No company on earth has a larger infrastructure solely for the purpose of tracking, identifying, and eliminating threats before they attack. At Symantec, we take cybersecurity very seriously. That's why we work 24/7 to protect your systems. Learn why 99% of the Fortune 500® depend on Symantec for their security.

SECURITY.

SYMANTEC.COM/EVERYWHERE

Confidence in a connected world.

symantec™

SECURITY RESOURCES

SOLUTION:

As the pace of information growth accelerates, so does the risk and threats to compromise that information. Attacks today are proving to be more sophisticated, well-organized, and covert in nature. Businesses now require a focus on security continuity more than ever—a focus that allows them to continuously respond to internal and external changes. Moving forward, businesses need to develop a security strategy that is risk-based and policy driven, information-centric, and operationalized across a well-managed infrastructure. Symantec can help by introducing security solutions designed to secure and manage IT assets:

- **Develop and Enforce IT Policies** – Control Compliance Suite 10.0 will deliver greater visibility into an organization's IT and compliance risks, to ultimately provide better intelligence at lower costs.
- **Protect Information** – Data Loss Prevention Suite 10.5 addresses key concerns about preventing data loss through social media and protecting information in private clouds.
- **Manage Systems** – Altiris IT Management Suite 7.0 offers complete IT management that enhances effectiveness through faster deployments and increased security, reduces costs by closing technology gaps, and improves manageability amidst the increasing information and infrastructure sprawl.
- **Protect Infrastructure** – Symantec Protection Suites will offer in-depth protection tailored to specific areas of the IT infrastructure and provide businesses with unified information security management across endpoints, servers and gateways

SOLUTION:

Thawte is a leading global Certification Authority. Our SSL and code signing digital certificates are used globally to secure servers, provide data encryption, authenticate users, protect privacy and assure online identifies through stringent authentication and verification processes. Our SSL certificates include Wildcard SSL Certificates, SGC SuperCerts and Extended Validation SSL Certificates. Founded by Mark Shuttleworth in South Africa, Thawte was the first certificate authority to issue SSL certificates to public entities outside of the United States, quickly accounting for 40 percent of the global SSL market. In 2000, Thawte was acquired by VeriSign, Inc. and has become a key member of the VeriSign family of trust brands. To date, Thawte has issued more than 945,000 SSL and code signing certificates since 1995, protecting identities and transactions in over 240 countries.

CONTENTS

John Savill (john@savilltech.com) is an MCITP: Enterprise Administrator for Windows Server 2008, VMware Certified Processional, and an 11-time MVP. His latest book is *The Complete Guide to Windows Server 2008* (Addison-Wesley), and you can check out his technology videos at http://www.ittv.net.

Microsoft Security Solutions

As you know, Microsoft has put a lot of effort into the security of all its products and has come a long way, with features such as BitLocker, Address Space Load Randomization, Microsoft Baseline Security Analyzer, IPSEC, AD Rights Management Services, and many more. But to be clear from the start, our focus will be on the solution sets available outside the core operating system, those service or infrastructure components that we should never neglect as valuable partners of Microsoft's powerful built-in solutions.

We'll start with those components that can help increase your security with no associated licensing costs (of course, the deployment of any solution will still have associated costs in the form of testing, deployment effort, communication, and so on). I should point out that these solutions run only on genuine copies of Windows; so you won't be able to use them if you're running unlicensed machines.

Malware Protection

Let's look first at the feature that is conspicuously absent in Microsoft's basic OSs: they include no built-in, comprehensive, real-time malware and antivirus protection in the box. Although the systems do offer a good, stateful firewall and Windows Defender spyware protection, Microsoft's only provision in this regard is part of standard Windows Update processes, when clients will download and execute a Malicious Software Removal Tool once a month (the frequency of the tools virus definition update). However, this tool is performing only a one-time check for a very specific list of the most damaging and prevalent Internet viruses. There is no other built-in protection between these scans, so the Malicious Software Removal Tool is not a substitute for, but only complements, a full-featured antivirus solution. (Note that you can run the Malicious Software Removal Tool at any time. Just go to http://www.microsoft.com/security/malwareremove/default.aspx and execute it.)

To address this lack, Microsoft released Microsoft Security Essentials (MSE), its free malware protection solution that includes scheduled scans and real-time protection. MSE is based on the same code as the enterprise Forefront malware protection solutions (more on those later) and is the replacement for the Windows Live OneCare consumer solution, which Microsoft has retired. MSE is a thin and light malware protection solution that typically goes unnoticed. It does use more resources during the scheduled complete system scans, but those usually run in the "quiet" hours. If you are not currently running an antivirus solution, you should definitely evaluate MSE as a first step toward securing your system (it's available for download at http://www.microsoft.com/security_essentials/).

I should note, however, that although MSE gives a good level of protection for home machines, it is not really recommended in a corporate environment where security and data sensitivity are as a rule more important than on a user's home computer. For corporate environments, the Microsoft solution is its Forefront Client Security, which provides more in-depth protection, including spyware and rootkit protection, than MSE offers. Also note that MSE is not manageable in an enterprise, and its End User License Agreement (EULA) does not allow enterprise use. In contrast, management is a key component of Forefront Client Security; all aspects of Forefront's deployment and management can be centrally configured through Group Policy and the Forefront Server Security Management Console, which offers automatic server discovery and web-based management.

Base protection of the core OS and file system, such as Forefront Client Security offers, is critical; however, some data stores and services require more customized solutions. For those needs, consider an Exchange database or SharePoint store because traditional antivirus solutions will not adequately protect the information in the stores or the type of transactions that are occurring—for example, opening and inspecting an ESE database and hooking into mail flow. Imagine what would happen when a typical antivirus solution identifies a virus in a database file. The solution would block access to the file, stopping the mail service, and then rip out the virus. In a database, this would represent corruption, so the solution requires intelligence.

To achieve this intelligent protection for specific workloads, other solutions are available as part of the Forefront suite; namely, Forefront Security for Exchange Server, Forefront Security for OCS, and Forefront Security for SharePoint. All of these solutions are tailored to the workloads and data flows

specific to the technologies, and each has various deployment options based on the level of protection and architecture used in each environment. Following are some of the key focus areas for security solutions with each of these applications.

- Forefront Security for Exchange Server—Provides protection for both antivirus and antispam, giving users a clean messaging environment. In addition to protection from malware and spam, Forefront Security for Exchange Server also supports keyword blocking, which provides additional protection from inappropriate or sensitive content. Forefront Security for Exchange Server supports the Edge, Hub, and Mailbox roles (not Unified Messaging or Client Access Server), and deployments generally fall into one of two architectures.
  The first deployment option provides ***baseline protection.*** In this scenario, Forefront is deployed only to the Hub and Edge servers, not to the Mailbox server. This may seem foolish if you consider that all the email resides on the Mailbox servers. However, when you consider the message flow of Exchange and understand that every piece of mail flows through a Hub transport server (e.g., even if a message is being sent to a recipient on the same Mailbox server as the sender), you can see that, even without protection on the Mailbox server, every piece of mail is still checked. Items that are not routed through the Hub, such as public folder content, sent items, and calendar data, however, are not protected.
  The second deployment option is known as ***global protection.*** In addition to running Forefront on the Edge and Hub servers in this deployment, you run Forefront on the Mailbox servers. And even in this global protection deployment, you are not wasting resources and delaying mail delivery by scanning a mail item multiple times.
  Forefront uses a stamp technology to mark messages that have been scanned with an AV stamp, so future hops on the mail route will not recheck. There is actually a third deployment type, which is a hybrid solution of Forefront Security for Exchange Server and the Microsoft-hosted protection solution, Forefront Online Protection for Exchange, (FOPE). I go into more detail later about FOPE.

- Forefront Security for SharePoint—Protects content in the SharePoint store and removes malicious code. Like the Forefront Security for Exchange Server solution, Forefront Security for SharePoint allows for blocking of content based on keywords and file filters. As an example, you could automatically block file types known to carry malware, such as executable files, or you could block content that can lead to legal exposure, such as mp3 files (Forefront protects us not only from malware but also from ourselves).
- Forefront Security for OCS—Like Forefront Security for SharePoint protection, Forefront Security for OCS provides file filtering and inappropriate content blocking in real time, and scans all content that passes through OCS for malware.

All three of these Forefront application security solutions use SmartScreen technologies, which in turn use up to five antivirus engines from major partners, with configuration options for how many engines should be used at any one time. Performance is an important consideration here. Turning on every antivirus engine may sound great; but doing so means every piece of data is being scanned multiple times, which will increase resource use and slow down communication. You also will receive diminishing returns as you use more engines. For example, one engine may miss a certain type of malware, and running three engines (the recommended number) is likely to offer comprehensive protection; but running five engines will almost double your overhead yet not give you much more protection than three engines. Antivirus engines available in Forefront include the Microsoft AV Engine, Norman Virus Control, Authentium Command Antivirus Engine, VirusBuster Antivirus Scan Engine, and Kaspersky Anti-Virus.

Forefront is pleasantly manageable, with comprehensive monitoring options that include a useful dashboard overview that lets you quickly check the health of your protection environment, provides a summary of work performed, and indicates the number of incidents; it also allows detailed investigation of incidents. As Figure 1 shows, you can see all

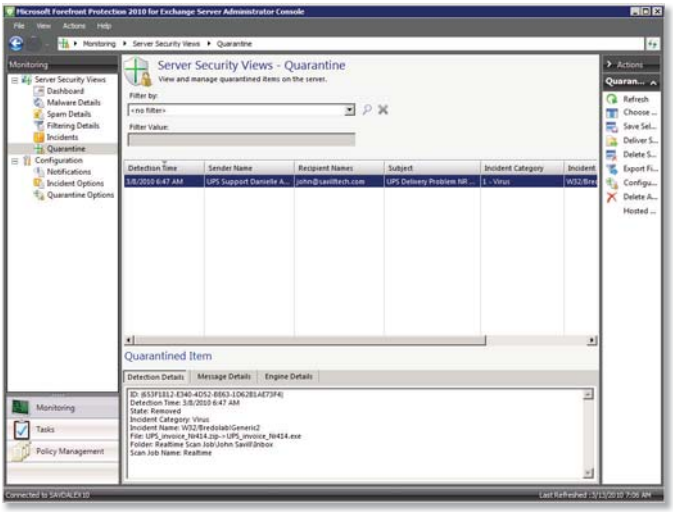

Figure 1: The Forefront quarantine view provides detailed information about quarantined items, including incident category and the engine that detected the item.

of the quarantined items, why they were quarantined, and which engine detected the viruses.

Last, but by no means least, when we consider malware protection we must include Forefront Online Protection for Exchange (FOPE). You may remember this service by its previous name of FrontBridge. FOPE is essentially a Microsoft-hosted antivirus, antispam, and disaster recovery (DR) solution that uses the same technologies that an organization can use locally. When I say DR solution, I'm not talking about a full replica of your messaging environment; there are other solutions for that, such as Business Productivity Online Standard Suite (BPOS). Instead, the DR protection is essentially providing store and forward for your messages. So if your email infrastructure is unavailable, senders will not receive nondelivery reports (NDRs); instead, FOPE will queue the messages and try to resend every 20 minutes for up to five days, which should give most organizations time enough to resolve whatever ails their messaging environment.

To use FOPE, you simply update your mail exchange (MX) DNS record to point to the FOPE environment instead of your own mail servers, and then configure the online protection service with a list of valid accounts in your environment. By not forwarding messages for which no valid recipient exists, this configuration provides an additional layer of optimization. You can do this account management manually, but most organizations opt to use the automatic synchronization functionality that is available. When FOPE receives a message, it does three things to protect your inbound traffic: It checks the message for malware using three different engines; it runs heuristics to ascertain a spam rating; and it compares against the valid recipients. It also can check any block/safe sender lists that end users manage. If the message passes all the checks, it is forwarded to the organization's own messaging infrastructure. If you also want FOPE to protect your outbound traffic, you can configure Forefront as your SMTP relay within the Exchange configuration.

Does this approach mean you don't run any internal protection? No; not at all. Using FOPE provides you with a great initial filter that will get rid of most all spam, malware, and invalid recipient data. But because you always want multiple layers of protection—defense in depth—you still need to run internal protection. The good news is that FOPE will filter out most of the junk so it never reaches your infrastructure, which saves your resources by allowing your infrastructure to focus on valid traffic. For internal traffic, however, an infected user system that is sending traffic to other internal users will never pass through FOPE and therefore will not be identified and cleaned. Figure 2 shows a typical messaging infrastructure, with the traffic flow. Note that this configuration gives comprehensive malware and spam protection, but it is not a complete security solution for the messaging infrastructure.

What I've described here is the standard spam and antimalware solution. There are additional cloud service options for Exchange, as well as encryption, archival, and business-continuity solutions that provide more than five days of mail retention and enhanced online email capabilities.

Batten Down the Hatches

Microsoft has a comprehensive set of malware-related security solutions; however, many other types of threats go beyond malware, such as network intrusion, hacking, and denial of service—the list is long. But all of these possibilities require access to your resources to pose a threat. So in addition to making sure that Windows Firewall is enabled on all your servers and clients, you need a first line of protection

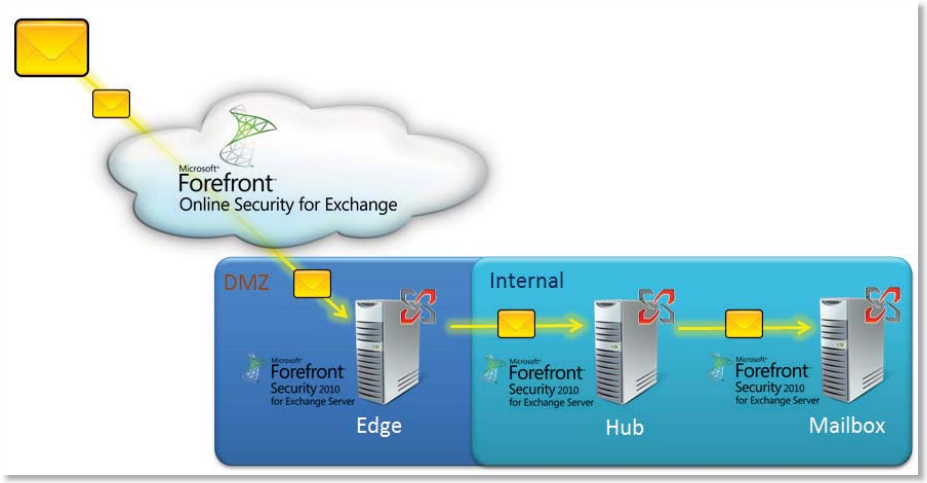

Figure 2: Forefront offers a complete spam and malware protection configuration for your Exchange environment, but it's not a complete security solution for the messaging infrastructure.

that checks data as it tries to enter and leave your network. You also want control over the data that's moving between your demilitarized zone (DMZ) and your internal network. The solution to both of these requirements is a firewall that performs deep-packet inspection of the data and, based on defined rules, allows the traffic through to specific targets or denies the traffic.

Many of you will be familiar with Microsoft Internet Security and Acceleration Server (ISA), which in the past was Microsoft's primary firewall solution. This product has evolved and has been rebranded Forefront Threat Management Gateway (TMG) 2010. TMG provides that critical network edge, gateway security access; however, thinking of TMG as just a firewall is like thinking of Microsoft Office as something that's pretty good for editing text files. In other words, TMG provides great firewall functionality, but it can do so much more.

First, TMG is stateful. Rather than just following rules independent of already-established communication, TMG's logic considers the entire traffic flow and initiation source of the connection to allow/block traffic.

Another heavily used feature is TMG's capability to publish internal services out to the Internet; this aspect allows easy access to key services from the Internet without production services being placed in a DMZ or Internet traffic being allowed through to the internal network. Commonly we see TMG publish Remote Access connectivity, Exchange OWA, and SharePoint, all of which are transparent to the end user. But we will discuss another, more comprehensive option, Forefront User Access Gateway (UAG), in the next section.

In addition to using TMG to secure traffic that is trying to get into your network, you can use it as the VPN connection points between a main data center and branch locations to help secure traffic between your locations. This use precludes the need for costly dedicated links or user-centric VPN configurations.

Come On In

If TMG is the burly security guard who is really there to provide protection from web-based threats, you can think of the next solution more as the friendly concierge who is there to accompany you and help you access the resources you need. Forefront UAG is an evolution of the Whale technologies that Microsoft purchased a number of years ago. UAG's primary focus is the secure and controlled enablement of remote access. UAG enforces very granular access controls and policies based on the services being accessed, the health of the machine that is connecting into the network via UAG, and the user credential.

The primary connectivity technologies in use today that UAG supports are SSL VPN, application publishing, and DirectAccess. As many organizations are examine their remote access strategy, they generally de-emphasize the VPN-style connectivity solutions because of the relatively complex usage for end users. In addition, the corporate IT team has no way to connect to user machines for maintenance. As you will see, DirectAccess provides a connectivity solution that is completely invisible to the end user, and it gives corporate IT access to user machines, which facilitates management such as patching and sending updates. If you are considering deploying DirectAccess, you'll also be looking at UAG. New with

Windows 2008 R2 and Windows 7 Enterprise and Ultimate editions, DirectAccess provides end users with a consistent connectivity experience to corporate resources, whether they're at work, at home, or at Starbucks.

With DirectAccess, users don't have to access VPN in the corporate environment before they access the resources; instead, DirectAccess behind the scenes connects to the target by encrypting the IPv6 traffic in IPSec. If IPv6 connectivity is not available point-to-point (such as over the Internet), the system will tunnel with 6to4 (if the client has a public address) or Teredo (if the client has a private NAT address). If that fails, the system will use IP-HTTPS (which tunnels IPv6 packets inside an IPv4 HTTPS session). But even with all these technologies, the target must support IPv6 or the connection will fail, and many organizations have substantial resources that are IPv4 only.

This is where UAG shines. UAG can act as a protocol gateway, converting the IPv6 communication into IPv4, thus allowing direct access to all resources in the organization, even if they don't support IPv6. This capability is a huge plus and will be key to many organizations' ability to adopt DirectAccess. UAG also can publish Remote Desktop Services applications and even expose certain internal network file system content.

Typically, you would purchase UAG as an appliance with UAG pre-installed. Alternatively, it is available as a Hyper-V virtual appliance, which is essentially a preconfigured virtual machine, and also as a software-only solution for Windows Server 2008 R2 (now 64-bit only). It's interesting to note that installing UAG also installs TMG, on which UAG is built. You can think of TMG as protecting your boundaries and checking outbound access, while UAG facilitates inbound access. Note that in this scenario you won't be able to use the full TMG functionality on your UAG box because it's not supported—TMG is there only for UAG to use.

Identity Management

Microsoft's identity management solution has been through more name changes than Prince. First known as Zoomit VIA before Microsoft's purchase, afterward it became Microsoft MetaDirectory Services (MMS). Later, it was called Microsoft Identity Integration Server (MIIS) and then Identity Lifecycle Manager (ILM). The new name, Forefront Identity Manager (FIM), brings the identity management solution into the Forefront brand.

Why do you need an identity management solution? For the answer, consider an average organization, which has ERP systems responsible for benefits, payroll, and performance management; directory services for authentication and authorization; messaging environments; and many other types of systems. Each member of the organization needs an account in all these systems, in addition to access to resources. All of this requires a significant amount of administrative effort; and when someone leaves the organization, each of those accounts must then be deprovisioned. An identity management solution has connections to all of the identity repositories within the organization; to facilitate synchronization of attributes between systems and provisioning actions, the solution allows relationships to be defined between objects in each system in a central metaverse. You can think of the metaverse as the convergence of all the data collected.

One common configuration for such a solution would be for the ERP system to be the authoritative account repository, and then the identity management system would pull information on new accounts from ERP through its management agent and automatically create accounts in the Active Directory (AD) and mailbox in Exchange. It would then add the users to groups based on their department and role, and would even execute scripts to create home areas. All this can all be done without any coding. Likewise, changes in one system, such as a change in job role or name change, can automatically be reflected in all the other systems. FIM also supports the synchronization of passwords between systems, including the AD, through a special agent that runs on each domain controller; this is a capability most other solutions do not have. If approval steps are required before certain actions are performed, then you can leverage the workflow capabilities of FIM to request approvals from specified people at certain points in the provisioning sequence.

As Microsoft's identity management solution, Forefront Identity Manager has a large number of connectors to support communication with systems such as PeopleSoft, AD, AD LDS, IBM Directory Server, Novell eDirectory, Oracle,

SQL, Informix, dBase, SAP, RACF, and Sun LDAP; in addition it supports various file formats such as CSV and DSML. I should also note that, for developers, FIM provides a framework that allows custom management agents to be written, which in turn permits FIM to connect to any system that is required.

A common source of effort for the Help desk of any company is dealing with countless user password reset requests, updating certain attributes, and handling group membership. FIM provides an end-user, web-based portal that enables self-service so that users can reset their own passwords and even update specified attributes of their identities. For group memberships, a Microsoft-provided add-in is available for Outlook so that if a user wants to join a group, he just submits a request. The FIM workflow automatically mails the group's owner for approval; once the approval is given (through Outlook or the portal), the user is joined to the group. The key here is that the end-user interaction is as intuitive as possible, so Help desk requests are decreased instead of increased as users struggle to use complex identity systems while FIM is painless and integrates with both their office tools and the web.

In addition to enabling the synchronization of identity information and provisioning capabilities, FIM provides management capabilities for your public key infrastructure (PKI); this includes allowing for automated configuration, provisioning, deprovisioning, and tracking through configurable auditing of your user digital certificates and smart cards. One improvement from ILM is that FIM now supports third-party certificate authorities through the opening of the FIM APIs.

Putting It All Together

As you can see in Figure 3, we actually combine all of the Forefront solutions for a complete security solution. This model has many layers of defense, starting with your network edge, which provides your first line of protection while it enables legitimate remote access through TMG and UAG. Your application services have protection specific to their workloads, and you have client security on your workstations and servers that works in conjunction with the OS's firewall. And the really good news is that access to most of these solutions is included for organizations with enterprise CALs, which makes implementation of Forefront a no-brainer.

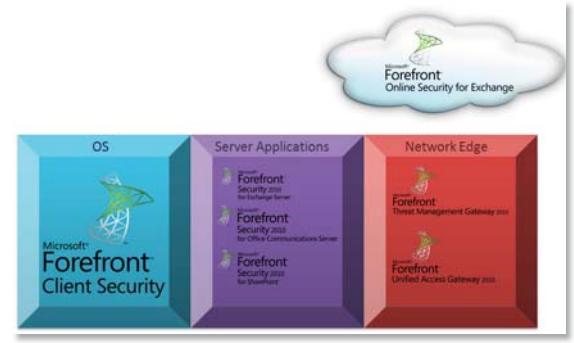

Figure 3: The complete Forefront security solution has many layers of defense.

Virtualization and Its Effects on Security

Virtualization has moved beyond cool to being a key technology organizations employ to save money and power and increase their responsiveness to ever-changing needs. In fact, 2009 was the crossover year in which virtual server installations overtook physical server installations. Where once an organization had 20 physical boxes, it now may have a single physical box that runs each OS instance in its own virtual environment. This consolidation results in multiple OS instances sharing hardware and introduces new challenges that you previously did not have to consider.

Hyper-V provides you with a secure-by-design hypervisor that ensures isolation between guest operating systems. This means that although virtual machines are running on the same physical piece of equipment, their interfaces do not allow communication, so they have no more access to each other than if they were physically separate boxes. The virtual machines must use the network to communicate.

In this context, security starts with the hypervisor, which is the thin piece of code running directly on the hardware that manages certain types of resource access, such as CPU and memory. This hypervisor runs no third-party code or drivers, which is important when you consider that it has direct access to the CPU and memory of every guest. So it's important to ensure that no malware runs within the hypervisor itself.

Most of the actual drivers run within the parent (also known as the root) partition, which is the OS you install on the physical hardware, and through which the Hyper-V role has been enabled. This configuration allows Hyper-V to take advantage of the existing large driver base available for the Windows OS. Within the parent partition, a virtualization stack performs the actual management of the guest partitions, in addition to handling intercepts and device emulation. A guest partition never talks directly to any hardware. Each guest has its own worker process in the parent, in addition to unique resource pools in the hypervisor; this arrangement helps to ensure isolation between guests and maintains your security.

As you can see in Figure 4, the Hyper-V model uses the VMBus for resource access to support communication between the guests and the parent. It is common to show this VMBus as a single bus shared by all the guests, but in fact each guest has a separate VMBus that links it to the parent to ensure isolation of guest traffic. Essentially, a partition is a security boundary, with nothing shared between partitions.

When we talk about security, we talk about trust—specifically, which components trust other components. In Hyper-V, all guests are untrusted by both the parent and the hypervisor itself, which means a guest that is running malicious code cannot inflict direct harm on another guest, the parent, or the hypervisor (of course, it can still try traditional over-the-network attacks and the like). Inversely, the parent partition is trusted by all guests and the hypervisor itself. This relationship stresses the importance of securing your Hyper-V parent partition because malware in the parent can affect everything else that's running on the box.

It's important to adhere to some best practices to help protect your environment as much as possible. First, you should think about your Hyper-V host and how best to lock it down. Remember that Windows Server Core is your friend. Introduced with Windows 2008, the Server Core installation option provides you with an operating system that contains a subset of the components of a full installation; for example, in 2008 R2 there is no Microsoft Management Console, no Windows Explorer, and no Explorer shell and subset of the .NET Framework. Basically, anything you don't need to run the services Server Core supports is not included. This economy is good

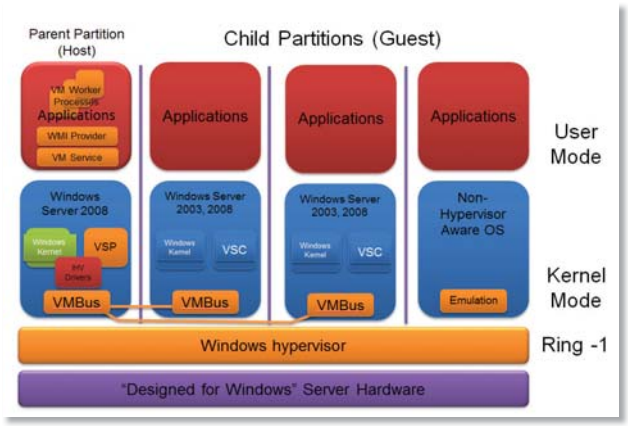

Figure 4: In this high-level picture of the Hyper-V architecture, note that each guest has a separate VMBus channel, so nothing is shared.

for security because fewer components means a smaller attack surface and fewer components to patch, which in turn means less downtime and management overhead.

You should use Server Core for your parent partition and make sure to keep it patched. Once you have your parent partition, keep its use limited to Hyper-V. If you have applications you need to run, run them in a guest partition. Keep the parent clean of any other software except the normal antivirus, firewall, and core management agents—don't install other roles, features, services, or applications. Keeping the parent clean helps, once again, to reduce your attack surface. Also make sure you follow KB 961804 (http://support.microsoft.com/kb/961804) when you use antivirus on your Hyper-V host because some specific locations should be excluded from scanning.

On the subject of patching, the standard best practice of making sure your environments are patched applies equally to your virtual environments. Make sure you have a process or solution in place to inform your team of new vulnerabilities and solutions, and how to test and deploy them as quickly as possible. Microsoft provides a number of ways by which you can be notified of security events (at http://technet.microsoft.com/en-us/security/dd252948.aspx). If you have virtual environments that have been turned off for a long time, patch them before you introduce them into your production environment. Do this by placing them in a quarantine network to perform the patching, or use an offline servicing tool such as Deployment Image Servicing

thawte®

Building Trust Around The Globe

When you want to establish trusted relationships with anyone, anywhere on the internet, turn to Thawte. Securing Web sites around the globe with:

- strong SSL encryption
- expansive browser support
- multi-lingual customer support
- recognized trust seal in 18 languages

Offering outstanding value, Thawte is for those who know technology. Secure your site today with a Thawte SSL Certificate.

www.thawte.com

and Management (DISM). System Center Virtual Machine Manager (SCVMM) has some very powerful features for patching virtual machines. Also make sure all your guests are running the latest version of the Hyper-V integration services. Doing this helps enforce time synchronization, which will make any audit log investigations easier because the host and guest times will be in sync.

Typically your servers have a certain security or trust level; for example, you may have high-security, medium-security and low-security servers. When you convert these servers to virtual machines, you need to maintain that same security; and since the guest is only as secure as the parent, you need to ensure that the Hyper-V host's security level always matches or exceeds the guest's highest security level. It is common practice to place guests with similar security requirements on the same hosts to simplify this security enforcement.

To help keep the parent clean, you need to limit who has rights to your parent partition, which means limiting who is an administrator. If people need permissions in the Hyper-V environment, grant them only the permissions they require, which you can do at a very granular level thanks to Hyper-V's use of Authorization Manager (AzMan). With AzMan, you can configure particular roles with actions such as start/stop virtual machines or configure devices, and then grant each role to particular users and groups. It is also possible to set a scope for these permissions so that different groups of users can manipulate only specific guests, rather than all guests on a host. Do not designate large numbers of people as administrators on the host; remember that the parent is a trusted source, and its security can affect the guest operating systems. If you are using SCVMM, an even richer delegation interface with three default defined profiles is available to you, to help you limit management rights. As an example, I could give the Sales Administrators VM management roles only for the Sales VMs. Figure 5 shows the main AzMan interface.

As you can see, I have created roles that encompass various permissions, and then I applied those roles to users and groups.

The next security items you typically need to consider are the virtual hard disks (VHDs) on the Hyper-V host. In the past, you would need to walk offsite with a 2U server stuffed under your shirt to steal a server's workload. A virtualized environment, however, consists of one or more virtual hard disks and some configuration files. If you can access the file system of a Hyper-V host, you can easily copy the virtual hard disks. Although such an act is not the theft of physical equipment, it may be far more damaging, given the data contained in the VHD—and a USB storage device is easy to conceal! If you steal the content of a Hyper-V host drive that contains 20 virtual machines, then you have stolen the equivalent of 20 servers. We will take measures to protect the Hyper-V hosts while we run them through normal Windows security, limiting administrative access and restricting NTFS permissions. However, anyone who has physical access to a server can reboot the box into an alternate OS and download information. To protect your hosts from this type of attack, you should use some type of volume-level encryption. A software solution would be to use BitLocker, which uses the Trusted Platform Module (TPM) on the motherboard to only allow access based on the presence of a password or USB device.

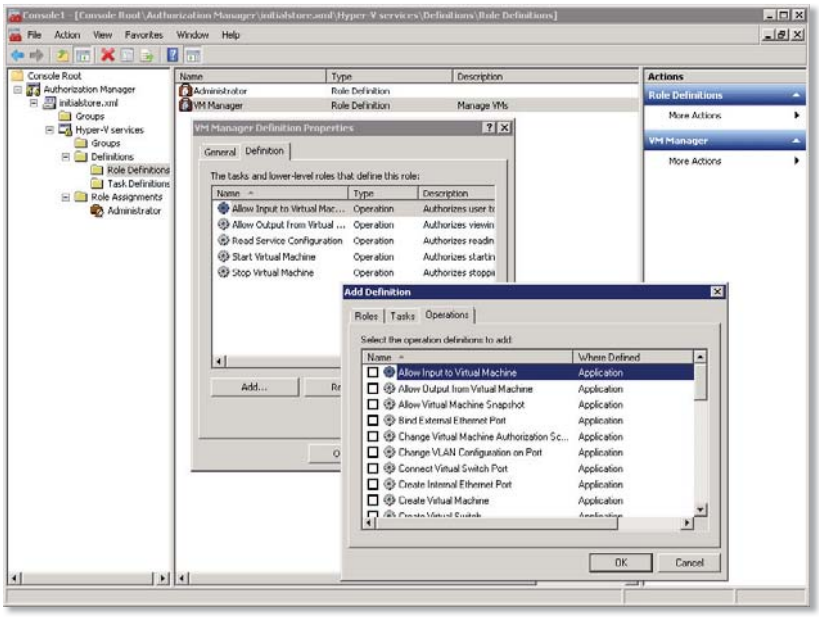

Figure 5: The main AZMan interface includes various operations that can be assigned to roles.

Hardware-based disk encryption solutions are also available. You need to ensure that the volumes that contain your virtual machines are protected, and a solution such as BitLocker has very little performance overhead associated with it—typically about 2 percent difference—when you compare it to an unencrypted drive. This encryption is transparent to your virtual machines, and without the password or USB key, it stops access to the volumes from outside the OS. A word of warning: If you do use BitLocker, make sure you keep the recovery key, which you need for access to the drive if the password/USB device is lost, in a secure location; without it, all that data will be unavailable. A best practice is to use Group Policy to configure BitLocker to store the recovery key in the AD as part of the host's computer object. When a virtual machine is no longer in use, make sure the virtual machine is backed up and then delete it. Once again, this approach reduces your attack surface by limiting the available information.

Networks introduce new challenges. For example, on the Hyper-V host, you want to limit Internet exposure as much as possible, so you always want to have a separate NIC just for the management of your Hyper-V host. And ideally, that NIC is connected to an internal management network that does not have direct Internet connectivity. You would then have one or more additional NICs configured as part of virtual networks, to be used by the guest partitions. The other advantages this configuration gives you are that if a problem occurs with the virtual switch, you still can connect through your management NIC to troubleshoot and resolve any issues, and your management traffic is not competing for bandwidth with the guest network workloads.

Now let's move on from separating the parent and guest traffic to isolating traffic between VMs on the same host. If you have guests that require connectivity to different physical networks, then the Hyper-V host will require multiple NICs that connect to each network, and then a virtual network linked to the NICs. As environments grow, this configuration becomes very cumbersome; so many organizations are adopting Virtual LAN (VLAN) tagging. Doing this allows each guest virtual network adapter to be configured with a specific VLAN tag, and only adapters with the same VLAN tag can communicate with each other. Alternatives would include leveraging technologies such as Internet Protocol Security (IPSec); however, IPSec has a greater list of requirements to implement than VLAN tagging does.

Protecting and Securing Internet-Facing Services

Just as no man is an island, no business is an island in today's world of Internet refrigerators with IP addresses. As your workforce become increasingly mobile, your organization faces unprecedented demands to offer services and easy access via the Internet to your customers, partners, and employees, even as you deal with increasing business security and regulatory requirements.

In the next section, we'll look at one primary scenario: how to best authenticate nonemployees for your Internet-facing services. There are numerous excellent solutions, such as DirectAccess, various <put your service name here> over HTTPS, RPC over HTTPS for Exchange, RDP encapsulated in HTTPS for remote access, and even VPN, for giving your employees access to the Internet. Many organizations are de-emphasizing VPN in favor of less user-intrusive technologies, but I have included it here because we are looking at the most secure ways to authenticate users.

Joe Public's Own AD Account

When you look at providing nonemployees with secure access to your organization's information over the Internet, your initial question should be "How do we authenticate these individuals and control what they can access?" The answer and solution will depend on the type of service they are accessing and the authentication mechanisms it supports.

In our perfect world, the service you are offering is to partner organizations, and it is web based. In this scenario, you can use Active Directory Federated Services (ADFS), a solution that leverages a requesting user's account in his local organization for access to resources in a partner organization.

A federated trust is created between the organizations out-of-band, which means there is no direct communication or link required between the organizations; the trust is set up by their sharing a certificate through snail-mail or perhaps email. Once organizations have established this federated trust with each other, their users can access resources in partner organizations without needing a separate account to do so. This approach has many advantages. For example, consider the normal provisioning of user accounts in partner organizations. You have to make requests to the partner organization to create an account; that takes time, and the user now has multiple accounts and passwords to remember, which often means numerous password reset requests. When a user leaves the organization, the process has to happen in reverse: You have to notify all your partners to delete these separate accounts to stop the user from continuing to access their resources. In contrast, with ADFS, the user accesses all the partner resources with his own account, and when his account is disabled, he also loses access to all partner resources.

Figure 6 summarizes how ADFS works. Essentially, ADFS servers in each organization have established the federated trust. When a user tries to access a resource in another organization, that user first communicates with the application server that is hosting the web-based service (1). The application server does not know the user, so it redirects him to the ADFS server in its local organization (2). The ADFS server does not know the user either, so it asks which organization having a federated trust the user is from (3). The user selects his organization and is redirected to its local ADFS server, which performs checks with the AD for details about the user. The local ADFS server then populates a Security Account Markup Language (SAML) token with claims about the user based on the AD content and digitally signs the token (4). The user is then redirected back to the ADFS server in the partner organization and the SAML token is presented (5). The partner ADFS server trusts the digital signature, consumes the claims made in the token, and creates a new SAML token for the user that it signs and then redirects the user back to the original application target (6). This time, the SAML token is presented, the application server trusts the local ADFS server token, and so it consumes the content of the token and grants access based on the claims made.

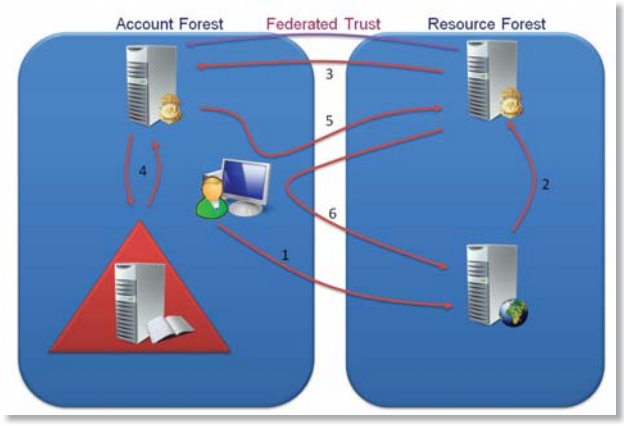

Figure 6: Here, using the ADFS solution, we see the redirections of the user's web browser.

Notice that everything in this sequence is a redirection of the user's web browser and occurs via the user's web session; there is never any direct communication between the organizations. A cookie is written to the user's computer so that the next time he tries to access the federated organization, he will automatically be redirected to his home federation server without manual interaction to select his home realm.

A clarification is in order here. Although I refer above to an ADFS server in each organization, that is not a requirement. Microsoft follows the industry-standard Security Assertion Markup Language (SAML) token specification. Version 1 of ADFS, however, uses the WS-Federation protocol sponsored by Microsoft, IBM, and VeriSign instead of the SAML protocol. The good news is that many vendors now support WS-Federation. This means ADFS can have a federated trust with any federation solutions such as Oracle SSO, Novell Access, Ping Identity Corporation SourceID, IBM Tivoli Manager, and others that support WS-Federation. And Microsoft supports most of the SAML protocol with version 2 of ADFS, which will increase the interoperability of ADFS and enable connectivity to more SAML systems.

Yet Another Account

So are ADFS and SAML the answer to all your organization's security issues? Sadly, no. Many organizations do not have federation environments today; and those that do may have only a limited set of federated trusts. ADFS simply will not work for typical individual customers as businesses move

beyond partner access. Even with today's ubiquitous IT availability, average users do not yet have their own AD and federation service. And if they do in another few years, managing all those federated trusts will be another challenge.

You therefore need to look at more "standard" solutions, which typically mean some kind of directory service. In the above scenario, you are offering services to the Internet, which means servers in the DMZ of the organization. Because the DMZ typically is far less trusted than your internal network, you want to be extremely careful about any infrastructure that is in the DMZ or that can receive communications from there. This means, for example, that you won't simply create a new organizational unit (OU) in your local AD called "Customers", create accounts, and then stick some domain controllers out in the DMZ— the exposure to your organization with this configuration would be huge.

Instead, a common solution is to use Active Directory Lightweight Directory Services (AD LDS), formally known as Active Directory Application Mode (ADAM). AD LDS provides a directory service that offers much of the functionality of the full Active Directory Domain Services (AD DS), including the capability to function as an authentication service; it also supports custom schema content. You can create an AD LDS instance in the DMZ that contains accounts for the accessing users; and if internal people also want access, you can use proxy objects in the AD LDS that refer to accounts in the main internal AD. The requirement for using AD LDS is that the application must support LDAP simple binds for user authentication because AD LDS does not support NTLM or Kerberos authentication. NTLM and Kerberos authentication are usually OK for many web portal application types, but they won't fit the bill for solutions such as SharePoint, which requires AD DS. Note that ADFS V1 can use AD LDS, but ADFS V2 will not do so out of the box, although it still can leverage AD LDS as an attribute store to help build the tokens.

So then what do you do if AD LDS won't work and you need AD out in the DMZ? Can you add an OU to your AD and place some DCs out in the DMZ? No. We essentially get into levels of comfort of what level of possible access our AD may have from the Internet. The less access the more comfortable we are.

The ideal solution would be to create an entirely separate AD forest for your DMZ use and potentially create a trust between the DMZ forest and your internal AD. Doing this would allow access for internal users to the resources offered in the DMZ. This firewall exception from the DMZ to your internal AD would allow communications only between the DCs in the DMZ and your internal network. Ideally, you also would encrypt this communication with IPSec. Although this configuration would be the most secure architecture, you can still go one step further to increase the security of your DMZ AD forest.

Windows 2008 introduced the concept of a read-only domain controller (RODC), which has a read-only copy of the AD database and is not allowed to make any changes or replicate changes to other domain controllers. Additionally, an RODC will not cache passwords for any accounts unless it is specifically set to do so. You can use these RODCs in the DMZ for your DMZ forest, have the writable DCs on the internal network, and open up specific ports in the firewall to allow the DCs to communicate. This way, servers in the DMZ can communicate with a local DC, which in turn will communicate with the full DCs on the internal network.

Another option would be to create a new domain in the same forest as the internal AD. However, to do this you need lots of communication between the domains because they are in the same forest. They would still share a common configuration and schema. Authenticated users would include all those people from the DMZ domain, and therefore this approach is discouraged in favor of an entirely separate AD forest. If you do use this extended corporate forest, make sure you place only RODCs, not full domain controllers, in the DMZ.

I should point out a caveat relative to the DMZ design in terms of using RODCs in the DMZ approach to AD: Some applications don't support RODCs. So when you are planning your DMZ AD design, make sure you understand the requirements of the applications you will be placing in the DMZ so you are confident that the RODCs will meet the requirements of the applications. For further information, Microsoft has a good guide (at http://technet.microsoft.com/en-us/library/dd728034(WS.10).aspx) that walks you through numerous models for AD in the DMZ.

High Availability/ Disaster Recovery for Virtual Environments

by Jason Bovberg

Ensure the business continuity of your imaginary technology

Virtualization is booming. But as increasing numbers of businesses incorporate the technology, it's important to remember that consolidating servers onto fewer physical machines comes with certain risks. Yes, businesses of any size can realize terrific resource benefits, but the cost of server failure can be high. A given host server can become quite valuable as it holds more and more virtual machines (VMs): If it goes down or experiences problems, business operations can be severely impacted. For this reason, you need a potent high availability/disaster recovery solution.

Microsoft and VMware offer such functionality within their industry-standard products, and those features provide good, basic protection. However, for ease of use or more granular functionality, you should consider third-party products such as those covered in this buyer's guide.

Virtualization Eases the Burden

To restore a physical environment, you have to laboriously re-install the OS and all applications; in a virtual environment, you're dealing with hardware-independent VMs—mere collections of files—that you can simply copy to alternative locations. Starting up the backed-up VM is a matter of dragging and dropping it to a newly operational host and starting it up.

Of course, virtualization in itself doesn't give you an automatic high availability/disaster recovery safety net. You still need an effective disaster-recovery plan for your VMs—just as effective a plan as you would have for physical machines. But virtualization does greatly simplify disaster recovery, particularly if you can—for example—utilize one of this buyer's guide's products to replicate your running VMs offsite, where they can await activation.

These products automatically transfer and restart your VMs at a predefined location; they can even start VMs in a certain order to proactively resolve any dependencies associated with the VM.

A prime method of ensuring high availability in a physical environment is failover clustering—and it's important in virtual environments, too. Failover clustering does a terrific job of increasing the availability of VMs—as well as the applications hosted inside those VMs—on the occasions of unplanned (and even planned) downtime. When the host server goes down, a standing-by host can automatically take its place, with working VMs ready to go. Microsoft's Hyper-V and VMware's VMware Server let you manage such cluster scenarios—for example, you can use VMotion or Live Migration to move and monitor workloads hosted in VMs.

Third-party tools can extend the protection offered by failover clustering across geographical boundaries, providing automated failover and failback. Primarily, these tools can provide geographical data protection, and of course they're just plain easier than implementing multisite clustering.

Considerations While Shopping

There are many considerations to keep in mind while checking out the market. Depending on your existing environment, the replication method is crucial. Will you use a SAN-based snapshot? Or will you use an image/snapshot of the VM directly? Another question to ask: Are all OSs treated equally? Will a Windows-based VM have an online snapshot backup taken, but a Solaris-based VM have to have a file-level or offline backup taken? What about support for mixed virtualization environments? You might have only VMware ESXi now, but you'd like to be prepared next year, when one of your users needs to use Hyper-V.

Most third-party tools simply add functionality to the basic feature sets available from Microsoft and VMware—or at least extend existing functionality. And in some cases, third-party tools might offer the same features but just do things better! Finally, watch for tools that offer smooth integration with suites or tool collections that you already have from the same vendor. If you're familiar with one company's tools, that familiarity will come in handy when using new tools.

A great product will provide that kind of integration with your existing infrastructure and capabilities to provide a comprehensive high availability/disaster recovery solution. As virtualization becomes more and more popular in the enterprise, high availability/disaster recovery capabilities are vital to business continuity.

And remember, disasters can come from many directions: Natural disasters, power loss, and simple human error can wreak equal havoc on your systems. So, hopefully, you're thoroughly testing your plan to make sure it works fluidly and effectively—and that advice holds true whether your environment is real or imaginary.

InstantDoc ID 125150

Jason Bovberg (jbovberg@windowsitpro.com) is a senior editor for *Windows IT Pro*, *SQL Server Magazine*, and *System iNEWS*, specializing in networking, hardware, storage/backup, and mobile and wireless. He has 20 years of experience as a writer and editor in magazine, book, and special-interest publishing.

Company Name	Toll-Free Number	Toll Number	Website	Product Name	Pricing/Licensing of Product	Virtualization Environments Supported	OSs Supported	Real-Time/Scheduled/ Batched Backups?
Acronis	877-669-9749	781-782-9000	www.acronis .com	Acronis Backup & Recovery 10 Virtual Edition	$1,999 per license	VMware, Hyper V, Xen, Parallels	Windows, Linux kernel 2.4.20 and above	Real-time/scheduled
BakBone	877-939-2663	858-450-9009	www.bakbone .com	NetVault: Backup VMware Plugin	Contact vendor	VMware	Windows, Linux	Scheduled
FalconStor Software	866-669-3252	631-773-5859	www.falconstor .com	FalconStor Network Storage Server (NSS)	Starts at $1,000 for the FalconStor NSS Virtual Appliance, plus $5,000 per managed TB; the FalconStor NSS Appliances start at $12,000 for 3TB of managed capacity after RAID.	VMware, Hyper V, Xen, Parallels	Windows, Linux, Solaris, AIX, HP-UX, Netware, and Macintosh.	Yes
Hitachi Data Systems	888-234-5601	408-970-1000	www.hds.com	Hitachi Data Protection and Disaster Recovery Solutions for VMware ; Hitachi Storage Cluster for Microsoft Hyper-V with Live Migration	Contact vendor	VMware vSphere and Site Recovery Manager, Hyper-V with Live Migration	Windows	Yes to all; depends on product and implementation
InMage Systems	800-646-3617	408-200-3840	www.inmage .com	InMage	$3,500 per server, $5,000 per appliance	All server virtual-ization platform products (e.g., VMware, Hyper-V, Xen, KVM)	Windows, Linux, Solaris, HP-UX, AIX	Continuous (real time), heterogeneous, async-hronous replication over IP with integrated WAN optimization
SteelEye Technology	800-769-4942	650-843-0655	www.steeleye .com	SteelEye DataKeeper & SteelEye DataKeeper Cluster Edition	$995-$3,000	Contact vendor	SteelEye DataKeeper will replicate any VM, regard-less of the OS, Hyper-V does have limitations.	Real-time, synchronous or asynchronous
STORServer	800-550-5121	719-266-8777	www.storserver .com	STORServer Agent for VMware Consolidated Backup	$45 per VM in quantity	VMware	Windows, Linux	Scheduled
Symantec	800-721-3934	541-335-7023	www.symantec .com	Storage Foundation for Windows High Availability	Starts at $4,495 per server	VMware, Hyper V	Windows	Real-time/continuous
Veeam Software		614-339-8200	www.veeam .com	Veeam Backup & Replication	$659 per CPU socket on source host; no license required for target host	VMware (vSphere 4 and VMware Infrastructure 3, ESX and ESXi)	Any guest OS supported by VMware	Yes: Near real-time replication, scheduled replication and/or backup, batched replication and/or backup

Editor's Note: Information in this buyer's guide comes from vendor representatives and resources and is meant to jump-start, not replace, your own research; also, it is not necessarily comprehensive, as some products might have been left out due to the writer's oversight.

	Automatic Failover/ Failback	Transport Mechanism	Distance Limitations	Support Multiple Replication Targets	Update DNS Entries Automatically?	Cloud Support	CDP Support	VM Image Restore, File-Level Restore, or Both?
	Instant Recovery (requires user intervention)	Contact vendor	None	Yes—dual destination, backs up in multiple locations	Contact vendor	Yes	No	Both
	No, but supports VMware failover using VCB	VCB/Network	None	N/A	No	No	N/A	Both as supported by VCB
	Yes	Download for the Virtual Appliance version. Shipping for Physical Appliances.	None	Yes	Can be scripted	Yes	Yes	Both
	Yes, Hitachi supports automated failover and data resynchronization	Fully integrated storage resource	Up to 150km non-disruptive; depends on product and implementation	Yes, including three- and four-datacenter configurations	Yes	Yes	Yes	Both
	Yes, failover and failback	IP	None (asynchronous)	Yes	Yes for both AD and DNS	Yes (already in distribution through various MSPs and cloud service providers)	Yes, true (not near) CDP with integrated bookmarking using application snapshot APIs (e.g., VSS, RMAN)	Data captured at block level, can restore volumes or files (depending on deployment model); also supports BMR for virtual servers
	Yes	TCP	None	Yes	Uses Windows Server Failover Clustering's "OR" Technology to facilitate failing applications across subnets	Yes. A mirror target can be a Windows Server instance on any cloud platform	No	Both
	No	Hotadd, SAN, LAN	None	Yes	Yes	Yes	No	Both
	Automatic failover	IP (for replication and basic network checks, but for our heartbeat for clustering is proprietary)	None—local, campus, global supported	Yes	Yes	Yes	Yes	Both, depending on configuration
	No, but includes tools to facilitate operator-driven failover and failback	TCP/IP or direct to local storage (for initial replica seeding or backup to removable media)	Network connectivity must be commensurate with the amount and frequency of data to be transferred	Yes, but not directly (can replicate to multiple targets individually)	Can be scripted	Yes, VMs can be replicated or backed up to the cloud	Supports near-CDP (replicate every 5-10 minutes)	Both: Can restore the entire VM or individual guest files and folders from Windows, Linux, Unix, Solaris, BSD, and Mac directly from a backup or replica

INSIGHTS FROM THE INDUSTRY

Top 10 Android Phones

When Android first arrived on the scene, it was disregarded by the enterprise as a consumer-friendly, hip platform for Linux geeks and 20 somethings who don't use phones for any real work. However, in just a year and a half, Android has transformed from a cool, open-source project to a very real contender in the mobility space.

A variety of events have triggered this change. First, Android garnered significant support from phone manufacturers such as Motorola and HTC, who have shifted their focus from Windows Mobile to Android. Second, platform-agnostic third-party vendors such as Zenprise have made supporting devices other than BlackBerry and Windows Mobile phones a much more realistic prospect. And third, I'd argue that the iPhone's popularity has made employees more bold (or more picky, depending on your opinion) in requesting phones they really want.

To see how ten of the hottest Android phones stack up, check out the buyer's guide table at www.windowsitpro.com, InstantDoc ID 125001. Below are some of the highlights for shoppers to consider.

Standard Features

Standard features that you can expect on any Android phone include:

- Camera and video
- Email (native Gmail support, and Outlook syncing through Exchange ActiveSync)
- GPS
- Wi-Fi
- Bluetooth
- Contacts management
- Touch screen/touch screen keyboard
- Android market access

Points of Differentiation

Despite all of these phones using the same OS, there are some significant points of differentiation to consider.

Exchange and Outlook support. All Android phones have ActiveSync, which allows for push synchronization between your Outlook account and your phone. However, many of these devices don't have native contact and calendar syncing, so if you're going to choose one of the devices that doesn't and you use Outlook, you'll need to download an app to sync them. The leading app for 2-way syncing is CompanionLink, which costs $39.99. Google also offers a free solution called Google Calendar Sync; however, you have to tie your Outlook account to a Gmail account in order for it to work, which will be an issue for some corporate accounts.

To see a table of 10 Android phones compared feature to feature, visit InstantDoc ID 125001.

Different Android versions. Each of the phones in this list either comes with version 1.5 (or 1.6) or 2.0 (or 2.1). Android 2.0 is a significant upgrade from the past version, but the only two Android smartphones that offer 2.0 are the Motorola Droid and the Google Nexus One. One of the most significant new features in 2.0 is contact syncing. See all the new features of Android 2.0 at developer.android.com/sdk/android-2.0-highlights.html.

Different carriers. Some individuals strongly prefer one carrier to another, and some organizations have corporate deals with a given carrier. As such, it's important to realize that many Android phones (and smartphones in general) only bundle with a specific carrier. If your carrier of choice is T-Mobile, then many devices are available. If you prefer one of the other three carriers, your options are more limited. The Google Nexus One offers the greatest selection, and is available on T-Mobile, AT&T, and Verizon.

Physical vs. virtual keyboard. If finger dexterity is your Achilles' thumb, you may prefer a physical keyboard, which would lead you to one of the sliders such as the Motorola Droid or CLIQ.

Best by Category

What device you use is a personal decision and will vary by individual, so I'm hesitant to make specific recommendations. Once you do decide which Android device you want (if any), I strongly recommend taking some time to see what users are saying across the web—while much of it might be inane, you should get some very good nuggets concerning the pros and cons from people that use the phone on a daily basis.

With that in mind, here is a quick list of the phone winners in each category (some categories, such as camera, I didn't factor because there are so many draws):

- Best processor: Google Nexus One
- Best memory/storage: Motorola Droid
- Best display size/resolution: Motorola Droid
- Best price: HTC Droid Eris, Motorola Backflip, and Samsung Moment
- Best battery life: HTC Hero
- Best variety in carrier coverage: Google Nexus One
- Lightest weight: T-Mobile MyTouch 3G

Be sure to check out the comparison table at InstantDoc ID 125001 for an in-depth look at how each device stacks up, and let me know your thoughts.

—Brian Reinholz

InstantDoc ID 125001

The 1&1 server totally configurable to your needs:

DYNAMIC CLOUD SERVER

A powerful virtual server environment with full root access. Adjust the processor core, RAM, and/or hard disk space to fit your needs. With the Dynamic Cloud Server, you can change your specifications at any time!

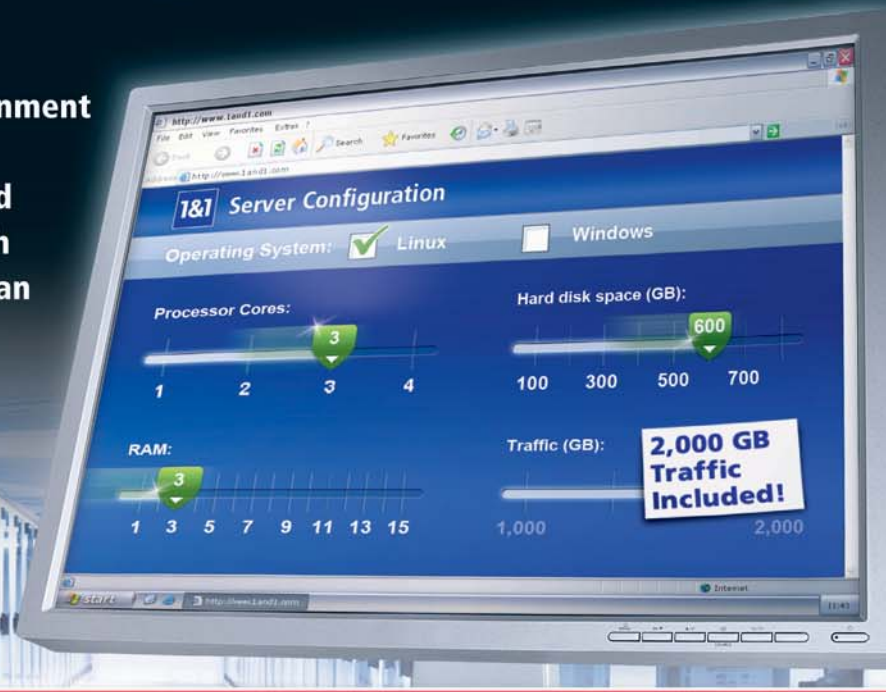

AMD Opteron™ 64

THIS COMPANY USES CERTIFIED bef CLEAN ENERGY

SPECIAL OFFER: 3 MONTHS FREE*

1&1® Dynamic Cloud Server – basic configuration includes:

- **1 Virtual Core of a Quad-Core AMD Opteron™ 2352 Processor**
- **1 GB RAM**
- **100 GB disk space**
- **Guaranteed resources** (just like a dedicated server!)

~~$49.99~~ per month

More server offers are available online. Visit our website for details.

*Offer valid as of May 1, 2010. Offer applies to Dynamic Cloud Servers only, up to a maximum discount of $149.97 per server. 12 month minimum contract term and setup fee apply. Prices valid for basic configuration only. For other configurations, additional costs apply. Visit www.1and1.com for full promotional offer details. Program and pricing specifications and availability subject to change without notice. 1&1 and the 1&1 logo are the trademarks of 1&1 Internet AG, all other trademarks are the property of their respective owners

MEMBER OF united internet

Call **1-877-GO-1AND1**

Visit us now **www.1and1.com**

1&1®

Capitals, Symbols, Numerals, and Whimsy

I once heard a person complain that developing the perfect complex password was the only way they could express their individual creativity in the homogenous, locked down, corporate soul crushing prison of the standard operating system environment. Having attained the right combination of capitals, symbols, numerals, and whimsy, they were loath to change it to something else after the three-week time limit had expired. Their reluctance was partly due to the effort required in remembering yet another complicated password, a task that can seem somewhat Sisyphean, and a fear that they would never again attain that level of creative ennui that resulted in their current password.

So how can we regularly create and remember complex passwords that need to be unintelligible enough not to be guessed by someone who knows us or replicated by someone who briefly catches a glimpse of us entering it through our keyboard? The answer is that we probably can't. Passwords are, to mangle a quote from Churchill, a terrible solution to the problem of authentication, except all the other ones we've tried.

Systems administrators are the worst when it comes to not changing their passwords. This is because they are the only people on the network who can ignore password policies and configure their accounts so that their passwords will not be changed.

When I was a new systems administrator at one workplace, I was approached by a contractor who asked me to reset his password rather than having to go through the lengthy process of password reset with the Help desk personnel. At that moment my phone rang and I asked the contractor to come back in a few moments. When he returned, he said not to worry about resetting the password because he had handled the matter. I asked him how he had resolved the problem. It turns out that the previous systems administrator had told him his password.

This sort of thing happens all the time. When I was making my transition from Help desk to systems administrator back in the 1990s, the guy who was our organization's current systems administrator went on leave for a few weeks and gave me the root password for a collection of servers. It was all the same password, but it had a cool mnemonic and I figured that he had changed it to that so I'd remember it and not bother him while he was on leave on some tropical beach somewhere. Fast forward five years where I've met up with that friend on an entirely different continent in one of the cages of a datacenter where there were servers he was responsible for managing. He told me that I'd be able to log on to them. I could.

—Orin Thomas

Security Recommendations for Microsoft's .NET Framework 4.0

Coinciding with today's release of the final version of Microsoft's .NET Framework 4.0, V.i. Laboratories makes certain recommendations regarding the security of applications developed on the framework. The company recommends that developers ensure their protection technologies offer the following:

- **Non-Invasive Protection** to ensure that intermediate language code is encrypted at the assembly level without impacting application code flow or introducing support dependencies.
- **Method Level Encryption** to prevent decompiling of .NET assemblies and minimize the exposure of decrypted application files or resources.
- **Anti-tamper Protection** to prevent tampering of specific application files or resources.
- **Protect APIs and SDKs** to allow the sharing of software IP with customers and partners without exposing sensitive code.
- **Support for ASP .NET** to protect web application code being hosted on third-party networks.
- **Platform Support** for compatibility with both 32-bit and 64-bit applications.

To learn more about V.i. Laboratories, visit www.vilabs.com.

—Lavon Peters

The latest 1&1 server solution for high performance needs:

HEXA-CORE TECHNOLOGY

NEW!

The ultimate in server technology, our powerful new hardware class is the perfect solution for running your resource-intensive applications.

THIS COMPANY USES CERTIFIED CLEAN ENERGY

SPECIAL OFFER: 3 MONTHS FREE*

1&1® Hexa-Core Servers – using the latest generation of AMD six-core processors:

- **2 x Six-Core AMD Opteron™ 2423 HE Processor**
- **Up to 32 GB memory**
- **Up to 2 TB of usable disk space with RAID 5**
- **Energy efficient, AMD-P technology**

Starting at ~~$499.99~~ per month

More server offers are available online. Visit our website for details.

*Offer valid as of May 1, 2010. 12 month minimum contract term and setup fee apply. Visit www.1and1.com for full promotional offer details. Program and pricing specifications and availability subject to change without notice. 1&1 and the 1&1 logo are the trademarks of 1&1 Internet AG, all other trademarks are the property of their respective owners. ©2010 Internet, Inc. All rights reserved.

MEMBER OF united internet

Call **1-877-GO-1AND1**

Visit us now **www.1and1.com**

1&1®

Process Automation

fueled by PowerShell®

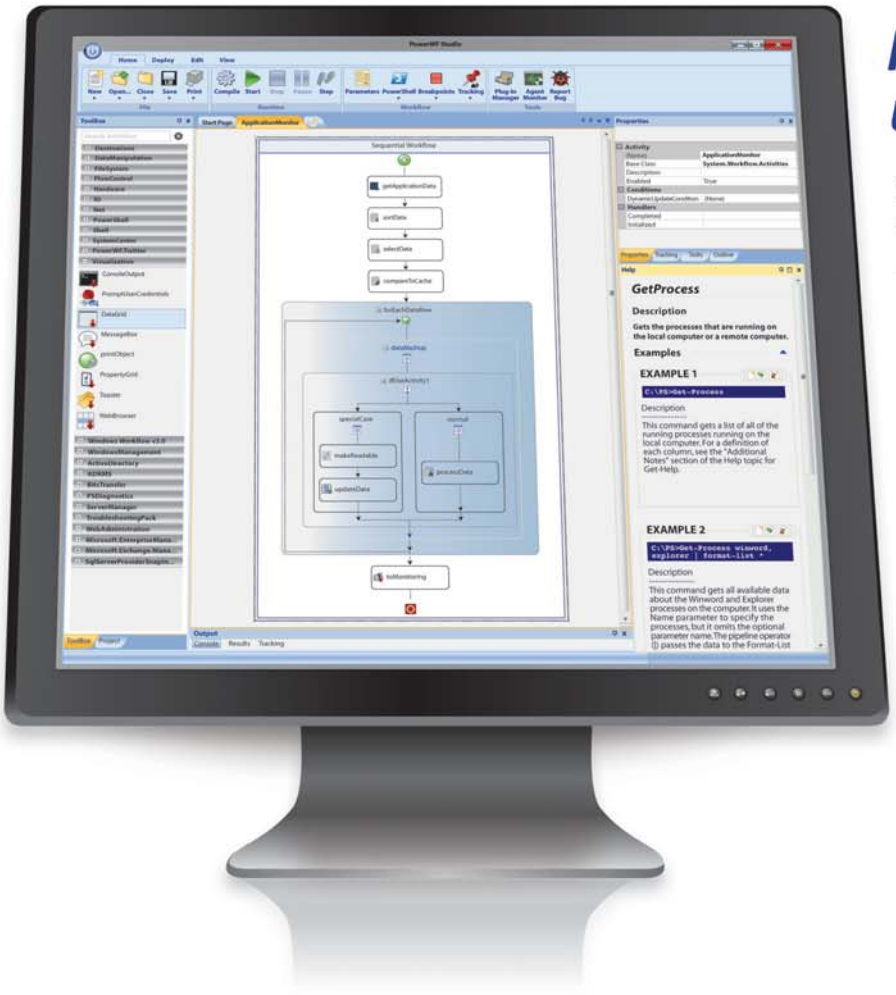

Now everyone can use PowerShell.

Discover every PowerShell module on a system

Visualize PowerShell as workflows

Extend with business logic

Integrate with a wide variety of 3rd party applications

Reuse workflows throughout the organization

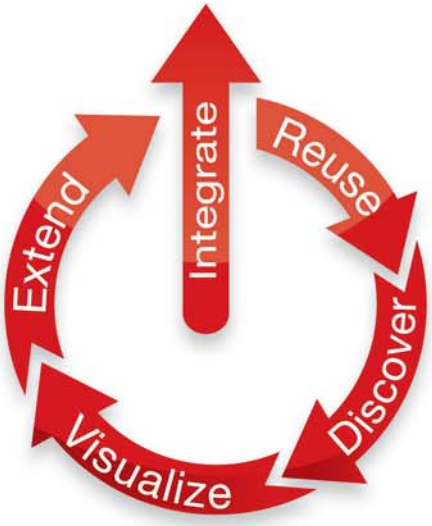

powerwf.com

314-590-5800 | 2008 Altom Ct., St. Louis 63146

sales@devfarm.com | Twitter: @PowerWF

PowerWF™ STUDIO

Visual PowerShell

Come visit us at Tech·Ed - Booth 2626.

PowerShell is a trademark of the Microsoft® Corporation.

For detailed information about products in this issue of *Windows IT Pro*, visit the web sites listed below.

VENDOR DIRECTORY

The following vendors or their products are mentioned in this issue of *Windows IT Pro* on the pages listed below.

DIRECTORY OF SERVICES | WINDOWS IT PRO NETWORK

Search our network of sites dedicated to hands-on technical information for IT professionals.
www.windowsitpro.com

Support

Join our discussion forums. Post your questions and get advice from authors, vendors, and other IT professionals.
www.windowsitpro.com/go/forums

News

Check out the current news and information about Microsoft Windows technologies.
www.windowsitpro.com/go/news

EMAIL NEWSLETTERS

Get free news, commentary, and tips delivered automatically to your desktop.

- *asp.netNOW*
- *DevProConnections UPDATE*
- *Exchange & Outlook UPDATE*
- *Security UPDATE*
- *SharepointPro Connections UPDATE*
- *SQL Server Magazine UPDATE*
- *Windows IT Pro UPDATE*
- *Windows Tips & Tricks UPDATE*
- *WinInfo Daily UPDATE*

www.windowsitpro.com/email

RELATED PRODUCTS

Custom Reprint Services

Order reprints of *Windows IT Pro* articles. Diane Madzelonka at Diane.madzelonka@penton.com.

Windows IT Pro VIP

Get exclusive access to over 40,000 articles and solutions on CD and via the Web. Includes FREE access to eBooks and archived eLearning events, plus a subscription to either Windows IT Pro or SQL Server Magazine.
www.windowsitpro.com/go/vipsub

SQL SERVER MAGAZINE

Explore the hottest new features of SQL Server, and discover practical tips and tools.
www.sqlmag.com

ASSOCIATED WEBSITES

DevProConnections

Discover up-to-the-minute expert insights, information on development for IT optimization, and solutions-focused articles at DevProConnections.com, where IT pros creatively and proactively drive business value through technology.
www.devproconnections.com

Office & SharePoint Pro

Dive into Microsoft Office and SharePoint content offered in specialized articles, member forums, expert tips, and Web seminars mentored by a community of peers and professionals.
www.officesharepointpro.com

NEW WAYS TO REACH *WINDOWS IT PRO* EDITORS:

LinkedIn: To check out the *Windows IT Pro* group on LinkedIn, sign in on the LinkedIn homepage (www.linkedin.com), select the Search Groups option from the pull-down menu, and use "Windows IT Pro" as your search term.

Facebook: We've created a page on Facebook for *Windows IT Pro*, which you can access at: http://tinyurl.com/d5bquf. Visit our Facebook page to read the latest reader comments, see links to our latest web content, browse our classic cover gallery, and participate in our Facebook discussion board.

Twitter: Visit the *Windows IT Pro* Twitter page at www.twitter.com/windowsitpro.

Windows IT Pro

by Jason Bovberg

Eat Me!

OUR favorite press release recently was a rather intricate April Fools joke, but we think the idea has some merit in our fast-food, environmentally conscious culture. On April 1, security company Intego announced that its packaging materials would be made from organically grown, edible materials. "Using rice- and soy-based packaging, and vegetable-based inks, all of Intego's boxes will be derived from organic raw materials, and will be edible." The company touted boxes with flavors such as bacon and roast chicken, with accompanying salad dressing—with, naturally, nutritional information printed on the sides. Watch for the "Eat Me!" logo!

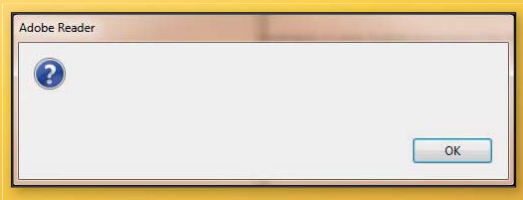

Figure 1: Yeah, I mean, really, WTF?

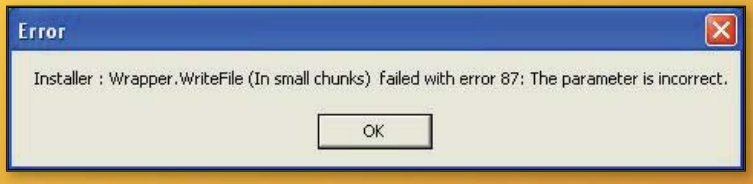

Figure 2: We prefer failures that occur in small chunks

5 More of Our Favorite Time-Wasting Flash Games

5 **Fly Hard**
www.kongregate.com/games/EmitterCritter/fly-hard

4 **First-Person Tetris**
www.firstpersontetris.com

3 **Open Doors 2**
www.kongregate.com/games/soapaintnice/open-doors-2

2 **Continuity**
www.continuitygame.com

1 **GemCraft Chapter Zero**
www.armorgames.com/play/3527/gemcraft-chapter-0

User Moment of the Month

I worked IT for an ad agency some years back. One day, I got a call from a user complaining that his screen was displaying nothing but shades of red. I guessed the problem must be related to an incorrectly connected monitor cable. I said, "Can you check the cable? Maybe it's not set correctly, and the blue and green pins aren't making contact...?" He insisted that the cable was firmly plugged in. Finally, I walked over to the user's office, finding immediately that the cable was incorrectly fitted, with bent pins. I showed the problem to the user, and he said, "That's not what I saw!" Then, the problem became clear: The user had been checking the other end of the cable. "Oh," he said, "I didn't know there were two ends."

—Jeremy

June 2010 issue no. 190, *Windows IT Pro* (ISSN 1552-3136) is published monthly. *Windows IT Pro,* 221 E. 29th St., Loveland, CO 80538, (800) 793-5697 or (970) 663-4700. Sales and Marketing Offices: 221 E. 29th St., Loveland, CO 80538. Advertising rates furnished upon request. Periodicals Class postage paid at Loveland, Colorado, and additional mailing offices. **POSTMASTER: Send address changes to *Windows IT Pro,* 221 E. 29th St., Loveland, CO 80538. SUBSCRIBERS:** Send all inquiries, payments, and address changes to *Windows IT Pro,* Circulation Department, 221 E. 29th St., Loveland, CO 80538. Printed in the USA.

Kiss your antivirus bloatware goodbye

NEW Premium Version with Firewall

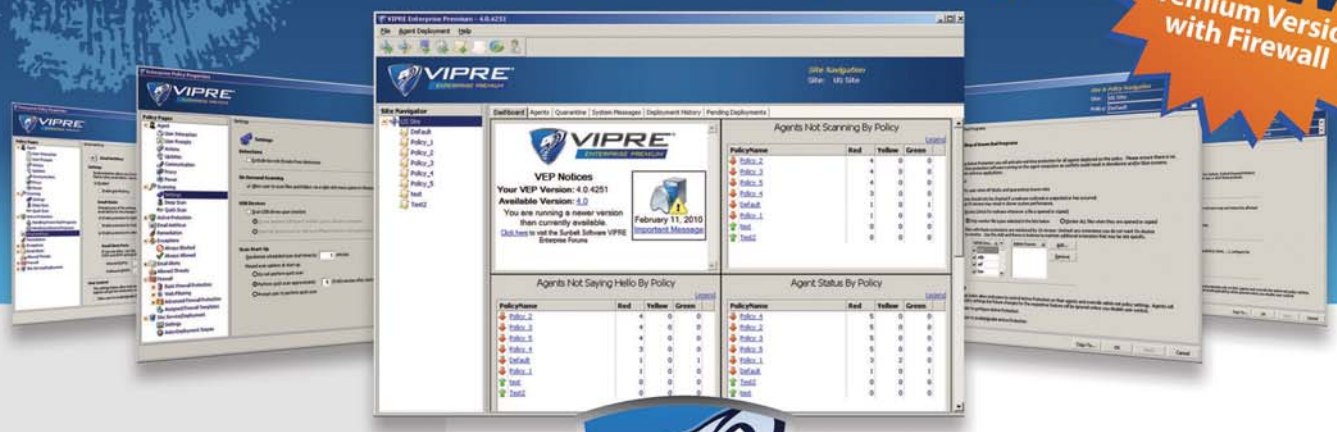

VIPRE®

TEST DRIVE ENTERPRISE PREMIUM

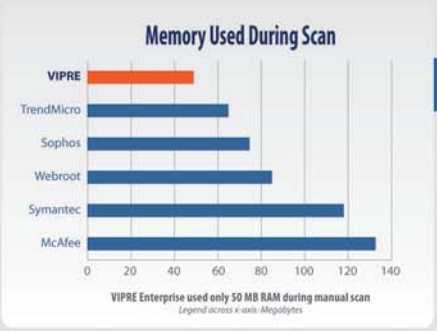

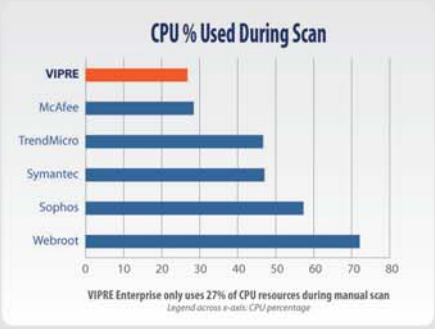

How does your current software compare? VIPRE Enterprise scans at a brisk 13.95 MB/sec and uses just 27% of CPU and 50 MB of RAM. In idle, it uses a mere 13.3 MB RAM with a disk footprint of just 113 MB. You'll hardly notice it's running!

Sunbelt Software

Special Competitive Upgrade: *50% Discount!*

Until now, antivirus engines have been Frankensteins, bolted together from bits and pieces of different products. They're slow, full of bugs, and hard to manage.

VIPRE Enterprise Premium is a revolutionary new approach. It combines high-performance antivirus, antispyware, and desktop firewall into a single agent so you get comprehensive endpoint malware protection with low system resource usage. It's fast, powerful and easy.

Plus, advanced anti-malware technology protects your system against the new wave of malware threats. No more juggling multiple programs. No more dealing with user complaints about slow workstation performance.

- **COMPLETE!** All-in-one protection from today's malware.
- **FAST!** High-performance and low impact on system resources.
- **EASY!** Manage everything easily from one command screen.
- **RELIABLE!** Configurable, real-time monitoring technology.
- **AFFORDABLE!** Ask for a quote with our 50% competitive upgrade discount!

Why struggle with slow resource hogs when you can manage ALL your malware threats with one fast, easy application?

Curious? Download your FREE copy of VIPRE Enterprise Premium and give it a test drive.

When you compare VIPRE Enterprise Premium to Symantec, McAfee, Trend Micro or whatever antivirus program you're using, **you WILL want to switch!** Don't worry, though. You can get VIPRE Enterprise Premium with a **50% competitive upgrade discount!**

Compatible with Windows 7

BEST IN TECH 2009 Computer Administrator

Windows IT Pro 2009 Community Choice GOLD

vb 100 VIRUS virusbtn.com

Plus we will buy out your existing maintenance contract for 1 year!

Download now: www.TestDriveVipre.com

Sunbelt Software Tel: 1-888-688-8457 or 1-727-562-0101 Fax: 1-727-562-5199 www.SunbeltSoftware.com sales@sunbeltsoftware.com

The IBM System x3550 M3 Express. When the downturn ends, the upside begins.

With new opportunities ahead, now is the time to invest in a faster, more powerful server: the IBM® System x3550 M3 Express® server, powered by the Intel® Xeon® processor 5600 series. By replacing your aging servers, the x3550 M3 can help you reduce operating costs, increase efficiency and respond to customers more quickly.

IBM System x3550 M3 Express

$3,299

or $84/month for 36 months[1]

PN: 7944E2U

1U dual-socket server featuring up to 2 Intel® Xeon® processor 5600 series

18 DIMM sockets 1333MHz DDR-3 (18 RDIMMs, 144GB max)

IBM System x3650 M3 Express

$3,065

or $78/month for 36 months[1]

PN: 7945E2U

2U dual-socket server featuring up to 2 Intel® Xeon® processor 5600 series

18 DIMM sockets 1333MHz DDR-3 (18 RDIMMs, 144GB max)

IBM System Storage DS3200 Express

$6,495

or $165/month for 36 months[1]

PN: 172622X

External Disk Storage with 3 Gbps Serial Attached SCSI (SAS) interface technology

Scalable up to 7.2TB of storage capacity with 600GB hot-swappable SAS disks

See for yourself.

See how much you could be saving – in just minutes – with the IBM Systems Consolidation Evaluation Tool.

ibm.com/systems/performance
1 866-872-3902
(mention 6N8AH27A)

intel Xeon inside™

Powerful. Intelligent.

[1]IBM Global Financing offerings are provided through IBM Credit LLC in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government customers. Monthly payments provided are for planning purposes only and may vary based on your credit and other factors. Lease offer provided is based on an FMV lease of 36 monthly payments. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice. IBM hardware products are manufactured from new parts or new and serviceable used parts. Regardless, our warranty terms apply. For a copy of applicable product warranties, visit http://www.ibm.com/servers/support/machine_warranties. IBM makes no representation or warranty regarding third-party products or services. IBM, the IBM logo, System Storage and System x are registered trademarks or trademarks of International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM trademarks, see www.ibm.com/legal/copytrade.shtml. Intel, the Intel logo, Xeon and Xeon Inside are trademarks of Intel Corporation in the U.S. and other countries. All other products may be trademarks or registered trademarks of their respective companies. All prices and savings estimates are subject to change without notice, may vary according to configuration, are based upon IBM's estimated retail selling prices as of 5/1/10 and may not include storage, hard drive, operating system or other features. Reseller prices and savings to end users may vary. Products are subject to availability. This document was developed for offerings in the United States. IBM may not offer the products, features, or services discussed in this document in other countries. Prices are subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or IBM Business Partner for the most current pricing in your geographic area.